

PIPELINE PROCESSING WITH AN ITERATIVE, CONTEXT-BASED DETECTION MODEL

T. Kværna, et al.

**NORSAR
PO Box 53
N-2027 Kjeller, Norway**

22 January 2016

Final Report

APPROVED FOR PUBLIC RELEASE; DISTRIBUTION IS UNLIMITED.



**AIR FORCE RESEARCH LABORATORY
Space Vehicles Directorate
3550 Aberdeen Ave SE
AIR FORCE MATERIEL COMMAND
KIRTLAND AIR FORCE BASE, NM 87117-5776**

DTIC COPY

NOTICE AND SIGNATURE PAGE

Using Government drawings, specifications, or other data included in this document for any purpose other than Government procurement does not in any way obligate the U.S. Government. The fact that the Government formulated or supplied the drawings, specifications, or other data does not license the holder or any other person or corporation; or convey any rights or permission to manufacture, use, or sell any patented invention that may relate to them.

This report was cleared for public release by the PRS OPSEC Office and is available to the general public, including foreign nationals. Copies may be obtained from the Defense Technical Information Center (DTIC) (<http://www.dtic.mil>).

AFRL-RV-PS-TR-2016-0080 HAS BEEN REVIEWED AND IS APPROVED FOR
PUBLICATION IN ACCORDANCE WITH ASSIGNED DISTRIBUTION STATEMENT.

//SIGNED//

Robert Raistrick
Project Manager, AFRL/RVBYE

//SIGNED//

Glenn M. Vaughan, Colonel, USAF
Chief, Battlespace Environment Division

This report is published in the interest of scientific and technical information exchange, and its publication does not constitute the Government's approval or disapproval of its ideas or findings.

REPORT DOCUMENTATION PAGE				Form Approved OMB No. 0704-0188	
Public reporting burden for this collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing this collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden to Department of Defense, Washington Headquarters Services, Directorate for Information Operations and Reports (0704-0188), 1215 Jefferson Davis Highway, Suite 1204, Arlington, VA 22202-4302. Respondents should be aware that notwithstanding any other provision of law, no person shall be subject to any penalty for failing to comply with a collection of information if it does not display a currently valid OMB control number. PLEASE DO NOT RETURN YOUR FORM TO THE ABOVE ADDRESS.					
1. REPORT DATE (DD-MM-YYYY) 22-01-2016		2. REPORT TYPE Final Report		3. DATES COVERED (From - To) 19 Mar 2013 – 22 Dec 2015	
4. TITLE AND SUBTITLE Pipeline Processing with an Iterative, Context-Based Detection Model				5a. CONTRACT NUMBER FA9453-13-C-0270	
				5b. GRANT NUMBER	
				5c. PROGRAM ELEMENT NUMBER 62601F	
6. AUTHOR(S) T. Kværna, D.B. Harris ¹ , S.J. Gibbons, and D. Dodge ²				5d. PROJECT NUMBER 1010	
				5e. TASK NUMBER PPM00018850	
				5f. WORK UNIT NUMBER EF122183	
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) NORSAR PO Box 53 N-2027 Kjeller, Norway ¹ Deschutes Signal Processing, LLC 81211 E. Wapinitia Road Maupin, OR 97037 ² Lawrence Livermore National Laboratory 7000 East Avenue Livermore, California 94550				8. PERFORMING ORGANIZATION REPORT NUMBER	
9. SPONSORING / MONITORING AGENCY NAME(S) AND ADDRESS(ES) Air Force Research Laboratory Space Vehicles Directorate 3550 Aberdeen Avenue SE Kirtland AFB, NM 87117-5776				10. SPONSOR/MONITOR'S ACRONYM(S) AFRL/RVBYE	
				11. SPONSOR/MONITOR'S REPORT NUMBER(S) AFRL-RV-PS-TR-2016-0080	
12. DISTRIBUTION / AVAILABILITY STATEMENT Approved for public release; distribution is unlimited. (OPS-16-10670 dtd 03 Mar 2016)					
13. SUPPLEMENTARY NOTES					
14. ABSTRACT Under existing detection pipelines, seismic event hypotheses are formed from a parametric description of the waveform data obtained from a single pass over the incoming data stream. The full potential of signal processing algorithms is not being exploited due to simplistic assumptions made about the background against which signals are being detected. A vast improvement in the available computational resources allows the possibility of more sensitive and more robust context-based detection pipelines which glean progressively more information from multiple passes over the data. Under this contract we have evaluated iterative components at different levels of the pipeline hierarchy. At the level of the raw waveforms, we have evaluated schemes for the detection and cancellation of noise transients which can both reduce the detection capability of a station and supply the parametric datastreams with phase detections which may result in spurious event hypotheses. At the single array level, the sensitivity to sites of monitoring interest can be diminished by energy arriving from other directions, for example from ongoing aftershock sequences. We demonstrate how energy from the direction of interest can be enhanced at the expense of the nuisance energy using adaptive beamforming and empirical matched field processing. The aftershock scenario can have significant consequences for the generation of near real-time bulletins both due to the increased number of events and the deterioration of fully automatic event bulletins due to spurious phase association. We provide proof-of-concept of a system for spawning a targeted process for accurate aftershock characterization in a given source region such that all associated phases are removed from the parametric datastreams. New iterations of Global Association, or equivalent algorithms, would then read in pre-screened detection lists and have a lower likelihood of generating false events. Empirical signal detectors for multiple repeating sources could eliminate associated signals from the parametric datastreams.					
15. SUBJECT TERMS aftershock sequences, repeating explosions, detection framework, pattern detectors, correlation detectors, subspace detectors, matched field detectors, nuclear explosion monitoring					
16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT Unlimited	18. NUMBER OF PAGES 122	19a. NAME OF RESPONSIBLE PERSON Robert Raistrick
a. REPORT Unclassified	b. ABSTRACT Unclassified	c. THIS PAGE Unclassified			19b. TELEPHONE NUMBER (include area code)

This page is intentionally left blank.

Table of Contents

List of Figures and Tables.....	iii
Acknowledgements.....	vii
Summary.....	1
1 Introduction	2
2 Methods, Assumptions, and Procedures.....	6
2.1 Cancellation of Repetitive Transients in Seismic Array Data Streams	7
2.1.1 Introduction.....	7
2.1.2 Mathematical Formulation.....	12
2.1.3 Algorithms	16
2.1.4 Spitsbergen Array Example	26
2.1.5 Summary.....	33
2.2 A Revised Detection Framework Architecture	35
2.2.1 Introduction.....	35
2.2.2 Test of the Signal Cancellation Component in the Detection Framework.....	37
2.2.3 Summary.....	51
2.3 Empirical Matched Field Processing and Adaptive Beamforming	52
2.3.1 Introduction.....	52
2.3.2 Application: Detection of Signals of Monitoring Interest During an Ongoing Aftershock Sequence	60
2.3.3 Application: Detection of Signals of Monitoring Interest in the Coda of a Large Earthquake	64
2.3.4 Application: Detection of Signals of Monitoring Interest Among Mining Explosions and Earthquakes.....	69
2.3.5 Summary.....	72
2.4 Evaluation of Event Hypotheses	73
2.4.1 Estimation of Detection Thresholds for Specific Source Regions.....	74
2.4.2 Evaluating Automatically Generated Event Hypotheses	78
2.5 Iterative Schemes for Automatic Event Bulletins: Coping With Aftershock Sequences.....	83
2.6 An Autoregressive and Grid Search Approach to Aftershock Classification	87

3	Results and Discussions	96
3.1	Cancellation of Repetitive Transients in Seismic Array Data Streams	96
3.2	A Revised Detection Framework Architecture	96
3.3	Empirical Matched Field Processing and Adaptive Beamforming	97
3.4	Evaluation of Event Hypotheses	97
3.5	Iterative Schemes for Automatic Event Bulletins: Coping with Aftershock Sequences	98
4	Conclusions	99
	References.....	101
	Appendix: Levinson-Durbin Algorithm for Solution of Block Toeplitz Matrix Equation	103
	Abbreviations	108

List of Figures and Tables

Figures

Figure 1:	Block diagram of a system component to remove local transients from a station's data stream.....	3
Figure 2:	High level sketch of possible future pipeline functions that would drive context-based signal processing algorithms.	5
Figure 3:	Location of the SPITS array in relation to Novaya Zemlya, the former Soviet nuclear test sites, and seismicity in the European Arctic.	8
Figure 4:	Mock helicorder plot for the central element vertical sensor of the SPITS array for 1 June 2011.	9
Figure 5:	Block diagram of the classical cancellation algorithm.	10
Figure 6:	Schematic depiction of transient signal cancellation using intermittent and continuous source models of the interference.	25
Figure 7:	The rose diagram (red) shows the azimuthal distribution of 195,000 Rg-type phases reported in the SPITS detection lists for 2011.	27
Figure 8:	Helicorder type record of one channel (SPA1 BHZ) of the SPITS array (left), and cancellation residuals for the continuous canceller algorithm (middle) and the detecting canceller algorithm (right).	28
Figure 9:	Composite of cancellation residuals for four iterations of the detection algorithm (top) and differences between the residuals from successive iterations (bottom).	29
Figure 10:	A 3-minute long data segment from the SPA0_BHZ channel.	30
Figure 11:	Cancellation options for the 9-element array of SPITS vertical sensors show that it is possible to suppress icequakes without distorting desired regional signals.	31
Figure 12:	Cancellation residuals for the full array of 21 vertical and horizontal elements of the SPITS array show that larger numbers of channels mitigate distortion.	32
Figure 13:	Details of cancellation residuals for the three algorithms around the regional waveform shown in Figures 11 and 12.	33
Figure 14:	Flow of data through the revised detection framework architecture.	36
Figure 15:	3 days of SPA0-BHZ data which is dominated by signals from nearby icequakes.	37
Figure 16:	(Top) 94 detections produced by detector 92532 and (bottom) 148 detections from detector 92541 produced during the first run of the framework.	38
Figure 17:	The 49 detections produced by the remaining detectors. The blue-shaded signals also appear to be icequake detections, but, perhaps because of interfering signals, were assigned to unique detectors.	39
Figure 18:	3 days of SPA0-BHZ data prior to cancellation (top) and after cancellation (bottom). The traces are plotted at the same scale.....	40

Figure 19: This figure shows how partial cancellation can cause changes in detection time for cases when multiple instances of the nuisance signal are present in close proximity.	42
Figure 20: Signals that were detected in the first run of the framework, but which were not detected in the run with cancelling turned on.	46
Figure 21: Instances in which the framework produced detections on incompletely cancelled icequake signals.	48
Figure 22: The three cases where a valid signal was not detected in run 365 because the signal onset was in the blackout interval for an icequake detection.	49
Figure 23: The two cases where a signal failed to be detected by correlation detectors because of the presence of a strong icequake signal.	50
Figure 24: The three cases where a valid earthquake signal buried within an icequake signal may have been detected.	51
Figure 25: Teleseismic paths from earthquakes in Myanmar to three North American arrays. The path length to ILAR (the nearest array) is about 8950 kilometers.	57
Figure 26: Waveforms of Myanmar calibration event (left) and target event (right), recorded at ILAR. Waveforms from nine stations (IL01-IL09) are shown.	58
Figure 27: Comparison of empirical matched field processing (black) and plane wave (red) power for the 2011 3/24 M_w 6.9 event.	59
Figure 28: Geographic configuration of the test with superimposed Wenchuan aftershocks and the 2006 DPRK nuclear test.	61
Figure 29: 2006 DPRK test superimposed at natural amplitude among aftershocks of the 2008 Wenchuan sequence.	62
Figure 30: Detection statistics for four matched field processors operating on data from the ASAR array show that empirical matched field processing with MVDR sidelobe suppression improves detectability of a weak signal among strong aftershocks from a different location.	63
Figure 31: Great circle paths from the DPRK nuclear test site (red) and the Nepal earthquake sequence (white) to the ILAR array in Alaska.	64
Figure 32: Scenario waveforms obtained by superposition of signals (S1) recorded at ILAR for the 12 May 2015 Nepal earthquake and the P wave observations (S2) also recorded at ILAR for the 6 January 2016 DPRK nuclear test.	65
Figure 33: Detailed plot of the superposition of the target waveform (S2) in coda of the Nepal earthquake (S1). The embedded explosion P wave is unremarkable, being similar in size to bursts of scattered energy from the earthquake.	66
Figure 34: Power traces (with square root scaling) for frequency-incoherent matched field processors.	67

Figure 35: Power traces (with square root scaling) for frequency-coherent matched field processors also demonstrate that adaptive processing is effective in suppressing unwanted interference, allowing the existence of a target signal to be inferred.....	68
Figure 36: Superposition of continuous stream data from station PD01 and the 2013 DPRK explosion.	69
Figure 37: Frequency-incoherent power traces (with square-root scaling) for conventional (top) and adaptive (bottom) matched field processors.	70
Figure 38: Amplitude envelope of a single raw data channel (top) compared to the power output (with square-root scaling) of the frequency-coherent matched field processor (bottom).	71
Figure 39: P-phase detection capability estimates for the ARCES array in 2°x2° bins. The red stippled curves denote distances of 20, 95 and 144 degrees from ARCES.	74
Figure 40: Predicted detection thresholds in units of magnitudes for the 44 IMS stations (both array and 3-component, primary and auxiliary) with best predicted detection capability for the North Korea nuclear test site.....	75
Figure 41: Predicted detection thresholds in units of magnitudes for the North Korea nuclear test site (assumed location 41.28°N, 129.07°E) as a function of epicentral distance for the 44 IMS stations displayed in Figure 3.....	76
Figure 42: Locations of the stations displayed in the previous two figures. The most sensitive stations for this source region have the station codes shaded.....	77
Figure 43: Extract from SL3 automatic event bulletin from the International Data Center (IDC) from November 5, 2005.	79
Figure 44: Extract from Reviewed Event Bulletin (REB) from the International Data Center (IDC).	80
Figure 45: Waveforms for selected stations (optimized for given phases from the Kashmir region) with arrivals from the events listed in Figure 44 labelled.	81
Figure 46: Source-scan for triggering WRA P-phase at time 2005-309:23.37.43.	82
Figure 47: Event location estimates from the fully automatic SL3 bulletin from the International Data Center (IDC) in Vienna for April 25 to April 30, 2015, and (b) event locations from the analyst-reviewed Reviewed Event Bulletin (REB) for the same time period.	83
Figure 48: The classical processing pipeline displayed schematically for the case of an aftershock sequence recorded on a global network. See text for details.	85
Figure 49: The scenario whereby the processing pipeline understands that an aftershock sequence is developing in a limited geographical region and spawns an offline process to detect and locate aftershocks accurately such that the detection parameters associated with well-characterized events can be removed from the parametric data stream provided to the phase association algorithm.	87
Figure 50: ARCES array beam aimed at teleseismic P-arrivals from the Nepal earthquake, together with the continuous AR-AIC function for this particular beam.	88

Figure 51: Fully automatic event location estimates obtained in the first 3 days of the Nepal sequence using an iterative grid search method based upon local maxima of the AR-AIC functions and pre-calculated traveltimes.	90
Figure 52: For each of the four array stations listed, the lowermost trace displays detections (SNR as a function of time) for which the velocity and azimuth constraints hold.	92
Figure 53: Traces optimized for the detection of regional P and regional S at 3 stations within 10 degrees of the October 8, 2005, Kashmir earthquake.....	93
Figure 54: Continuous AR-AIC functions for the traces displayed in Figure 53.	94
Figure 55: Close-up of the traces displayed in Figure 54 in original form (upper panel) and in reduced form (lower panel).	95

Tables

Table 1: Parameters controlling detector creation and cancellation behavior as they were set for this experiment.	40
Table 2: The 17 detections from RUNID 365 that match (nearly) in time with icequake detections from RUNID 364.	41
Table 3: Accounting of all detections in RUNID 364 and RUNID 365 excepting 242 known icequake detections from RUNID 364 detectors 92532 and 92541 and 17 icequake detections from RUNID 365 detectors 92584 and 92594	43

Acknowledgements

Seismic data from many stations of the International Monitoring System (IMS) was obtained from the International Data Center (IDC) of the Comprehensive Nuclear-Test-Ban Treaty Organization (CTBTO) in Vienna.

This page is intentionally left blank.

Summary

Under existing detection pipelines, seismic event hypotheses are formed from a parametric description of the waveform data obtained from a single pass over the incoming data stream. The full potential of signal processing algorithms is not being exploited due to simplistic assumptions made about the background against which signals are being detected. A vast improvement in the available computational resources allows the possibility of more sensitive and more robust context-based detection pipelines which glean progressively more information from multiple passes over the data. Under this contract we have evaluated iterative components at different levels of the pipeline hierarchy. At the level of the raw waveforms, we have evaluated schemes for the detection and cancellation of noise transients which can both reduce the detection capability of a station and supply the parametric datastreams with phase detections which may result in spurious event hypotheses. At the single array level, the sensitivity to sites of monitoring interest can be diminished by energy arriving from other directions, for example from ongoing aftershock sequences. We demonstrate how energy from the direction of interest can be enhanced at the expense of the nuisance energy using adaptive beamforming and empirical matched field processing. The aftershock scenario can have significant consequences for the generation of near real-time bulletins both due to the increased number of events and the deterioration of fully automatic event bulletins due to spurious phase association. We provide proof-of-concept of a system for spawning a targeted process for accurate aftershock characterization in a given source region such that all associated phases are removed from the parametric datastreams. New iterations of Global Association, or equivalent algorithms, would then read in pre-screened detection lists and have a lower likelihood of generating false events. Empirical signal detectors for multiple repeating sources could eliminate associated signals from the parametric datastreams.

1 Introduction

Seismological data centers attempt to provide accurate and rapid characterization of global seismicity. This is true of the International Data Center (IDC) in Vienna, a part of the verification regime for the Comprehensive Nuclear-Test-Ban Treaty (CTBT), and also at National Data Centers such as the US NDC. The processing pipeline currently employed at such data centers dates back to the 1980s and 1990s and consists of a transformation of continuous waveform data from multiple stations into streams of parametric information describing discrete detected arrivals. The limited computing resources available at that time demanded that critical event-building functions (association, location and identification) are performed subsequently on the parametric data, since very intensive signal processing operations were computationally prohibitive. Expensive signal processing operations were minimized in such systems by performing only a single pass over the waveform data.

In such an architecture, only very simple assumptions are possible about ambient background conditions in the selection and development of signal processing algorithms. Key assumptions are that events occur in isolation and that the background consists of white noise, usually uncorrelated and of uniform power among array sensors. These assumptions are clearly false in important and not-infrequent circumstances. Pipelines do have procedures to disentangle overlapped and interleaved seismic arrivals, but these usually are confined to association algorithms operating on the parametric representation of detected arrivals. Such procedures also fail demonstrably in cases such as extensive aftershock sequences. After being reduced to its parametric description, waveform data is generally not revisited in the processing pipeline - except to allow analysts to re-examine or relocate events, deleting false events or adding new events as necessary.

The single-pass architecture imposes significant limitations on the signal processing algorithms used for detection and phase characterization. Among these is the inability to adapt algorithms to suppress or cancel interfering signals with known characteristics. Future pipelines could allow waveform data to be revisited by detection and characterization algorithms that exploit some understanding or model of the context in which they operate. We consider how a context-based processing pipeline could operate with iterative elements at different levels of the hierarchy. For example, at a single station, the incoming waveforms may contain repetitive transient signals from nuisance sources; examples include the generation of icequake signals in the glaciers surrounding the SPITS array, or Rg phases generated by industrial activity near the KSRS array. We envisage central data collection systems with autonomous components to characterize local sources of interfering signals near each station and remove or suppress those

signals from the data stream before it is passed to the main pipeline for detection and subsequent processing. Such a subsystem would generalize existing masking functions that handle dropouts and data spikes to handle interfering, propagating transients from local nuisance sources. Such a signal-cancellation system is displayed in Figure 1. Further along in the processing pipeline, we may be attempting to associate phases globally from an extensive aftershock sequence in a limited geographical region. If a semi-autonomous sub-process were able to characterize this sequence accurately using a specially targeted signal processing strategy, we may be able to exclude all of the associated phases from the parametric datastreams that the association algorithms take as input - reducing the likelihood of false associations.

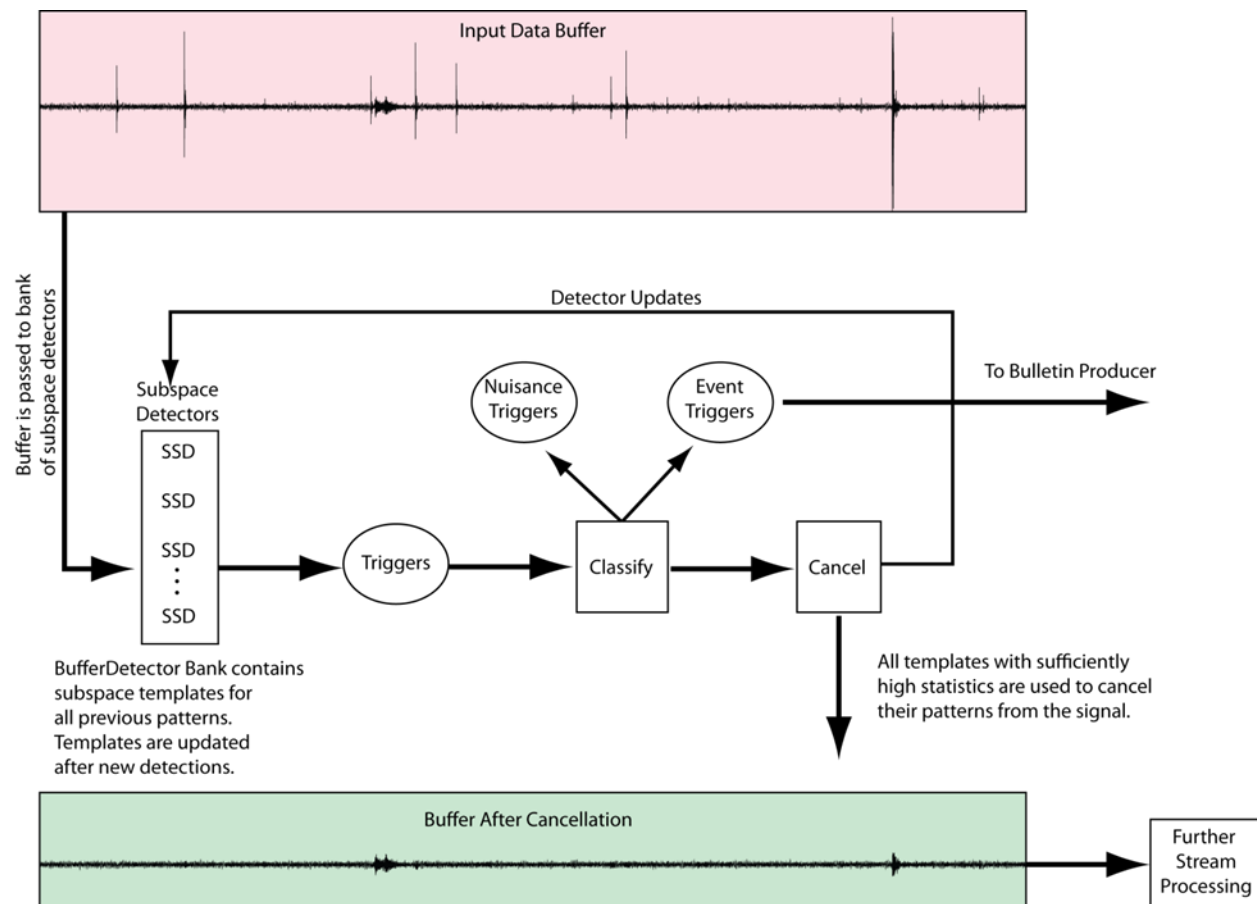


Figure 1: Block diagram of a system component to remove local transients from a station's data stream.

Another assumption about future pipelines is that they would maintain constantly-refined models (i.e. world views) of seismic activity around the globe, and use those models to adapt the signal processing functions of the pipeline to optimize detection and characterization functions. In such architectures, model state might be determined initially by pipeline processes not very different from those used currently. However, waveform data could be extensively

reprocessed with algorithms conditioned on hypotheses about the current state of seismic activity (e.g. large aftershock sequence in progress, mid-day mining explosions anticipated, etc.). A mature picture of activity may emerge after several iterations of model-driven signal processing on the continuous data stream. It is not that this picture of activity is of direct monitoring interest, but rather that it may allow more effective screening of interference in the continuous data stream.

Context-driven signal processing operations might include beamforming algorithms that search for events of interest among strong transient interference. Current beamforming algorithms are designed for very benign, white, uncorrelated noise background conditions. They are not well suited to detection among events in an aftershock sequence, for example. However, it should be possible to configure beamforming and related (for example, matched field processing) operations in second passes over data to make use of knowledge of interfering events gleaned from an initial pass. In this proposal, we give an example of such an algorithm: a data-adaptive, empirical matched field processor that substantially suppresses aftershock signals while remaining sensitive to signals from a known test site. The capability of this algorithm to suppress interference is significantly beyond the ability of a conventional beamforming algorithm because it is specifically designed to reject events with known characteristics. We here test this and related algorithms on a substantial fraction of a network with data recorded during an aftershock sequence.

The model-building supervisory function of future pipeline architectures also could instigate signal processing operations to support or refute hypotheses being formed or tested. For example, having formed a marginally-supported event hypothesis with a minimum of phase detections, the supervisor could direct special beams to search for phases predicted to exist at stations without detections. Current architectures have some capability to perform this function. For example, the IDC system performs associations, forms and locates events based on early-arriving seismic and hydroacoustic data (forming the SEL1 bulletin). The system automatically requests data from auxiliary stations based upon event hypotheses reported in that bulletin, and performs detections and measurements on those new data. However, this function can be substantially enhanced to direct re-examination of all waveform data, both to search for missing phases as noted and to confirm detections in the context of an hypothesized event (with location and magnitude estimates). An overview of the proposed pipeline architecture is displayed in Figure 2.

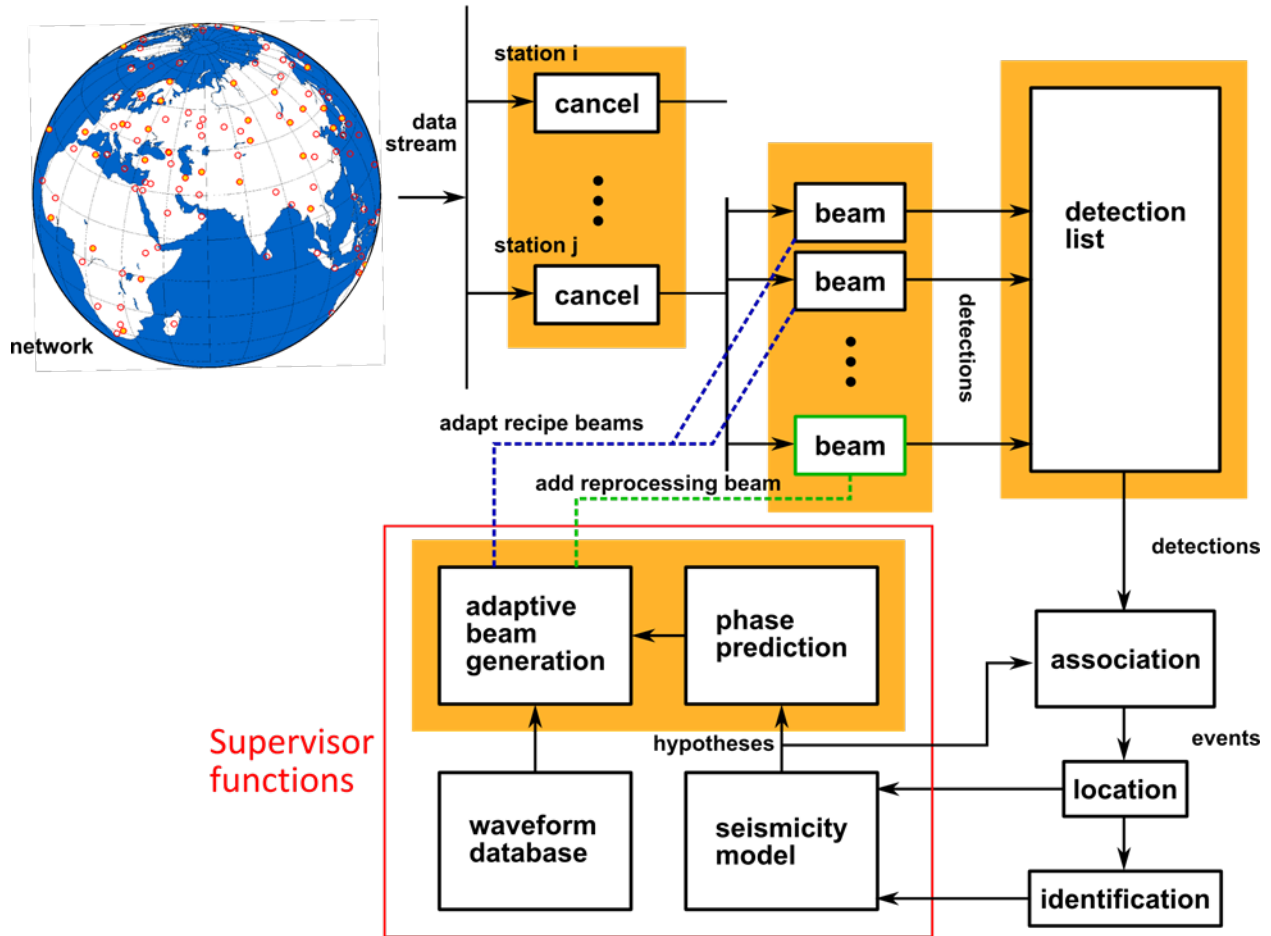


Figure 2: High level sketch of possible future pipeline functions that would drive context-based signal processing algorithms. Functions highlighted with an orange backdrop are ones that are prototyped as part of this project.

This report describes possible iterative approaches to several components of the processing pipeline. Each is considered as a separate subtopic under the Methods and Results sections.

2 Methods, Assumptions, and Procedures

Each of the following subsections describes a study of how a component of a modern processing pipeline could be implemented to provide an iterative approach to event detection and location.

Section 2.1 describes an evaluation of 2 different approaches for the removal of high amplitude nuisance transients from the incoming datastreams: a discrete, detection, signal cancellor and a continuous signal cancellation algorithm. The problem is highlighted for a case study of icequakes near to the SPITS array on Svalbard.

Section 2.2 examines the effects that the implementation of such a cancellor can have on operations of the detection framework.

Section 2.3 describes the application of empirical matched field processing and adaptive beamforming as a form for contextual signal processing to improve detection and association functions without the need for the “white noise” background assumptions on which current pipelines are run.

Section 2.4 examines the possibilities for evaluating the validity of event hypotheses with the use of targeted signal processing. Under current pipeline processes, association algorithms only have streams of parametric data available. The event hypotheses generated are currently evaluated only by an analyst, and we take a look at the steps that could be taken in an automatic system to eliminate false event hypotheses prior to analyst review.

Section 2.5 considers the scenario of an extensive aftershock sequence which generates a rapid succession of seismic phases globally which overwhelms existing phase association algorithms. We consider a solution whereby a separate offline sub-process is spawned to do a source-region specific search of the aftershock zone. With as many as possible of the largest aftershocks well-located, we can then remove all associated phases from the incoming data-streams.

2.1 Cancellation of Repetitive Transients in Seismic Array Data Streams

Highly repetitive local seismic interference often complicates operations at key stations in earthquake and verification monitoring networks. In some applications it may be desirable to remove interference transients from the continuous data streams of seismic arrays as a preprocessing step before normal detection, phase identification and event formation functions. We present several algorithms for transient interference suppression based on empirical models for interference sources. The general approach these embody is a dual of the traditional Widrow-Hoff noise canceller. Cancellers assume that a separate recording of an interference source time history is available and that this time history is correlated with components of unwanted noise in a primary data stream. The objective of traditional cancellers is to estimate a transfer function – in seismic applications, a Green’s function – relating the continuous source to interference in the primary stream, allowing the interference to be estimated and removed from the stream. By contrast, our application assumes that the Green’s function is known – through past observations of repetitive transients – leaving the source time history to be estimated. This is a more ambitious, but feasible, cancellation operation. We examine two algorithms, one which makes a continuous estimate of the interference source time history and a second which models the source as producing a relatively sparse set of discrete, impulsive events. In the latter case, correlation detectors find these discrete excitations to build an impulse-train representation for the source.

We give an example of icequake cancellation at the SPITS array in Svalbard and find that the continuous-source estimator is highly effective at removing interference transients at the cost of distorting desired signals. Distortion is reduced as the number of array channels is increased, but would require an unrealistic number of channels for high-fidelity preservation of desired signals. The intermittent source algorithm largely avoids distortion of desired signals but is not completely effective in removing interference. It is sufficiently effective, as we demonstrate, for significant improvements to pipeline operations. We also developed a post-processing correction to the continuous-source canceller that appears to be very effective in mitigating the distortion problem. The corrected continuous canceller would be the algorithm of choice for this application, were it not for its very high computational cost.

2.1.1 Introduction

Often it is the case that key seismic stations, otherwise ideally placed for some monitoring function, are located close to sources of seismic interference. This observation is true, in particular, for stations located in polar regions, in which the movement of glaciers and sea ice can generate thousands of small “icequake” events in a day. These events can cause major

difficulties for monitoring pipelines, particularly in the detection and association phases of operation. A particularly clear example is seen in the SPITS array located in the European arctic (Figure 3). SPITS is a key station for monitoring the Arctic region, including the former nuclear test sites on Novaya Zemlya. For events in the Novaya Zemlya region, SPITS has, during normal background noise conditions, the capability to detect events down to magnitude 2.0 [Gibbons et al., 2011].

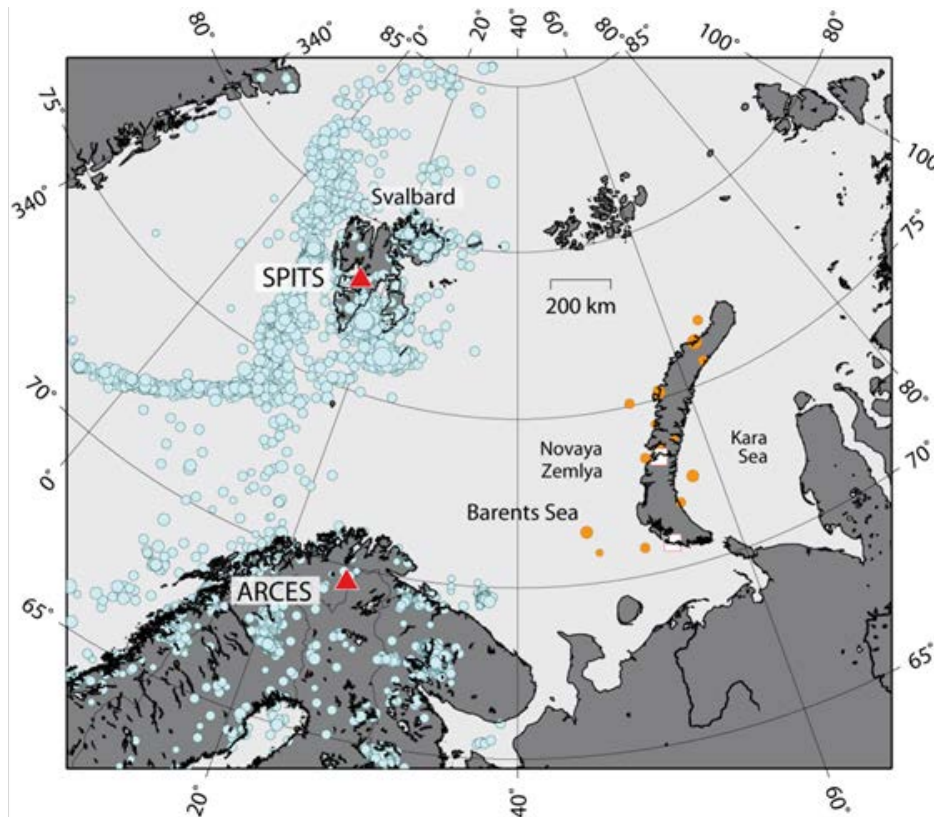


Figure 3. Location of the SPITS array in relation to Novaya Zemlya, the former Soviet nuclear test sites, and seismicity in the European Arctic. The blue symbols indicate location estimates of seismic events from the NORSAR-reviewed regional event bulletin between 1998 and 2008. The red/white squares indicate the approximate locations of north and south Soviet nuclear test sites on Novaya Zemlya. The orange symbols refer to events on or close to Novaya Zemlya in the time period 1992-2010. Symbol size relates to event magnitude.

However, the area around the SPITS array is characterized by permafrost and several neighboring glaciers which regularly generate seismic disturbances. During certain time intervals, e.g. in melting or freezing periods, these transients can dominate the seismic data and significantly reduce the detection capability of the seismic array. An example is shown in Figure 4 in the form of a helicorder type plot for 1 June 2011. The time interval between 15:00 and 18:00 UTC contains a large number of such transients.




Figure 4. Mock helicorder plot for the central element vertical sensor of the SPITS array for 1 June 2011. (See <http://www.norsardata.no/cgi-bin/spdatashow.cgi?year=2011&doy=152&sta=SPI>). Signals from events at regional distances have wavetrains up to several minutes in duration. The vertical scale is dominated by very short duration signals (duration 5-10 seconds) of which many thousand are detected in this 24 hour period.

Frequently undesirable local transient waveforms are repetitive, being produced by a relatively small number of sources of recurrent events. When this situation obtains, a pre-processing strategy for “tagging” or removal of nuisance transients based on waveform correlation techniques is feasible. This strategy could take the form of detection and marking for transients, much like glitch or dropout detection and marking is implemented with a masking channel in

some systems. A more ambitious approach might remove nuisance transients by waveform estimation and subtraction. This paper examines this latter possibility.

Waveform estimation and subtraction closely resembles noise cancellation approaches used to remove continuous interference (e.g. noise from rotating machines); we choose to frame the problem in a mathematical structure close this formulation (Figure 5).




Figure 5. Block diagram of the classical cancellation algorithm.

Noise cancellers, such as the adaptive Widrow-Hoff [1960] algorithm, model the observed waveform as a sum of a desired signal, s , undesired interference, y , and possibly uncorrelated background noise, η . The methods require an independent observation, r , of the interference process which is free of the signal one wishes to observe. This independent reference must be correlated with the interference in the primary signal channel. Cancellers, then, estimate a finite-duration transfer function, w , to “shape” the reference interference measurement to match (as \hat{y}) the interference in the primary channel. The estimate \hat{y} is subtracted from the primary waveform to suppress y . This kind of interference suppression can be highly effective and simple to apply, provided it is possible to identify and instrument the sources of interference with separate, dedicated sensors [for a geophysical application, see e.g. Harris et al., 1991]. It is possible to apply cancellation against multiple sources of interference

simultaneously with a suite of reference observations targeting all of the sources. This extension of the process leads to a multichannel variant of the simple canceller shown in Figure 5.

In the application we contemplate in this paper, an independent observation of the noise source is not available, but an estimate of the transfer function is known, being measured from past observations. Our application is the more difficult dual of the common cancellation problem, in that we propose to swap the roles of r and w , estimating an interference source of effectively unbounded duration.

We consider two approaches to the problem of estimating the interference source. In the first we treat the source as a completely general waveform, solving a least-square problem for estimating the discretized samples of the interference source. This approach, though completely general, has two principal disadvantages: the first is that its computational cost is huge, and the second is that it biases estimates of the desired signal. The second approach makes a simplifying assumption that the interference consists of highly repetitive transient waveforms, so that the source consists of a series of delta functions and the transfer function is identical to the repetitive waveform that we seek to remove from the data. This approach first performs a detection step, using a generalized correlation detector to identify the number and times of the impulsive events that constitute the interference source. A cancellation step follows detection, scaling the impulses to match amplitudes of predicted and observed occurrences of the interfering transients.

In fact, our model for the interference transfer function is somewhat more nuanced. We generalize the transfer function with a subspace representation to allow a range of variation in signals being targeted for cancellation. In principle, this generalization allows some variation in source location, mechanism or time history [Harris, 1989]. Although we have described the source time history as largely quiescent but punctuated by bursts of impulsive activity, in fact the source may be characterized by a near continuous range of activity with occasional very large excitations subsiding to a range of smaller excitations grading to an indistinct rumble. Because the detection step requires that we choose a threshold for detection, the smaller occurrences of interfering transients will not be removed from the data. The advantages of this approach are that it is comparatively much faster than the first method and, for desired signals, largely distortion free. Its disadvantage is that it removes less interference, particularly low-level interference. We observe that as the threshold is reduced the second method asymptotically approaches the first.

We also describe a post-processing technique for reducing bias in desired signals in the first cancellation method, based on a heuristic for identifying when the method is attempting to cancel a transient that is not a copy of the interference waveform.

This application and its solutions have much in common with finite fault decomposition (source slip estimation) using empirical Green's functions. The difference is that we consider problems with just a few source locations and continuous streams, rather than finite duration waveforms decomposed to hundreds of source locations gridding a fault surface. In addition our objective is not directly estimation of the source, which is responsible for unwanted interference in our application, but removal of the consequent transient interference in the continuous stream data.

The data we consider are multichannel: array data. In fact, the number of available channels controls a tradeoff between desired removal of transient interference and undesired distortion of waveforms from events of interest. The larger the number of data channels, the more reliably interference can be suppressed while simultaneously minimizing distortion of desired waveforms.

The paper is organized as follows: in the next section (2.1.2) we present our mathematical model of transient interference intermixed with waveforms from events of interest. In section 2.1.3, we derive the two principal solutions to the source estimation problem, discussing practical algorithms and the heuristic for suppressing bias in the continuous-source cancellation method. Following that, in section 2.1.4 we apply the methods to data from the Spitsbergen array, comparing results.

2.1.2 Mathematical Formulation

We assume the seismic observations are multichannel, with N_c channels, and represent the sampled continuous signals as a vector sequence:

$$\mathbf{x}[n] = \begin{bmatrix} x_1[n] \\ x_2[n] \\ \vdots \\ x_{N_c}[n] \end{bmatrix} = \mathbf{s}[n] + \mathbf{y}[n] + \boldsymbol{\eta}[n] \quad (1)$$

Streams from the individual channels are designated by $x_i[n]$; $i = 1, \dots, N_c$ with time index n . We use italic characters to indicate scalar quantities and bold characters to indicate vectors or matrices. The observations consist of three components: unmodeled waveforms from events of interest, $\mathbf{s}[n]$, undesired interference transients, $\mathbf{y}[n]$, and background noise, $\boldsymbol{\eta}[n]$. We assume each of possibly several interference components to have the form:

$$\sum_k \mathbf{g}[n-k]\mathbf{f}[k] \quad (2)$$

The $\mathbf{g}[n]$ are finite duration transfer functions (Green's functions or sums of Green's functions) relating forcing functions $\mathbf{f}[n]$ at the interference sources to the resulting transients observed in the data. N_T is the duration of the transfer functions in seconds. In equation (2) $\mathbf{g}[\cdot]$ is assumed to be $\mathbf{0}$ for values of its index outside of the interval $[0, \dots, N_T - 1]$. Taking into account the possibility of N_S independent interference sources, the complete representation for the data sequence is:

$$\mathbf{x}[n] = \mathbf{s}[n] + \sum_{i=1}^{N_S} \sum_k \mathbf{g}_i[n-k]\mathbf{f}_i[k] + \boldsymbol{\eta}[n] \quad (3)$$

As mentioned in the introduction, interfering transients often are repetitive. We use this fact to form empirical estimates of the transfer functions for groups of related interference events. Following [Harris, 2006], we use a 5 step process:

1. Construct a pool of events of all by running a power (STA/LTA) detector over a period of data when interference sources are active. Extract event data segments of duration N_T .
2. Cross-correlate waveforms from all events in the pool.
3. Cluster the events using a hierarchical agglomerative algorithm using waveform correlation as a distance measure. The clustering threshold determines the number of clusters. Select clusters which have more than some minimum number of events. The number of such clusters is an estimate of the number of discrete sources producing interfering transients. We expect such sources to be prolific; in fact, it is this characteristic which distinguishes nuisance transients from events of interest.
4. For each significant cluster, align the corresponding waveforms using delays estimated from peaks in the cross-correlation functions, then construct a data matrix of aligned waveforms with one column for each event.

5. For each significant cluster, construct an orthonormal representation for the column space by performing a singular value decomposition (SVD) of the data matrix and retaining the left singular vectors corresponding to the top singular values of the decomposition. We determine the dimension d (rank) of the representation by finding the d largest singular values with cumulative energy greater than some specified fraction of the total energy in all singular values. The matrix of left singular vectors is orthonormal by construction and constitutes the representation of the transfer function for a particular source.

Step 4 requires some elaboration. Since the data are multichannel, we multiplex the waveforms for a single event in channel-sequential form to form a column of the data matrix. This produces a consistent packed representation for a vector waveform (referring to equation 1):

$$\mathbf{x} = \begin{bmatrix} \mathbf{x}[0] \\ \mathbf{x}[1] \\ \vdots \\ \mathbf{x}[N_T - 1] \end{bmatrix} \quad (4)$$

that preserves moveout among the waveforms across an array. The dimension of \mathbf{x} is

$N_G \cdot N_T \times 1$. With N_E events, the complete data matrix has the form:

$$\mathbf{X} = [\mathbf{x}_1 \quad \cdots \quad \mathbf{x}_{N_E}] \quad (5)$$

It is common to normalize the event data so that $\mathbf{x}_i^T \mathbf{x}_i = 1$ for all events, so that no single event or subset of events comes to dominate the data matrix and the corresponding SVD.

The SVD of step 5 produces the decomposition:

$$\mathbf{X} = \mathbf{Z} \mathbf{\Sigma} \mathbf{V}^T \quad (6)$$

with diagonal singular value matrix and orthonormal left \mathbf{Z} and right \mathbf{V} matrices of singular vectors:

$$\mathbf{\Sigma} = \begin{bmatrix} \sigma_1 & 0 & \cdots & 0 \\ 0 & \sigma_2 & \cdots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \cdots & \sigma_{N_E} \end{bmatrix} \quad (7)$$

We select the dimension of the representation d as the smallest integer that satisfies:

$$\frac{\sum_{i=1}^d \sigma_i^2}{\sum_{i=1}^{N_E} \sigma_i^2} \geq \gamma \quad (8)$$

where γ is called the fractional energy capture threshold. Here the singular values are assumed to be sorted into descending order. The partition of \mathbf{W} that we retain as a representation for the collection of related waveforms is:

$$\mathbf{U} = [\mathbf{z}_1 \quad \mathbf{z}_2 \quad \cdots \quad \mathbf{z}_d] \quad (9)$$

The dimension of \mathbf{U} is $N_c \cdot N_T \times d$ and it is an orthonormal matrix:

$$\mathbf{U}^T \mathbf{U} = \mathbf{I}_{d \times d} \quad (10)$$

Referring to (2), our working hypothesis is that \mathbf{U} spans the matrix of Green's functions:

$$\mathbf{G} = \begin{bmatrix} \mathbf{g}[0] \\ \mathbf{g}[1] \\ \vdots \\ \mathbf{g}[N_T - 1] \end{bmatrix} \quad (11)$$

so that $(\mathbf{I} - \mathbf{U} \mathbf{U}^T) \mathbf{G} \approx \mathbf{0}$.

The representation matrix \mathbf{U} can be partitioned into separate time steps:

$$\mathbf{U} = \begin{bmatrix} \mathbf{U}[0] \\ \mathbf{U}[1] \\ \vdots \\ \mathbf{U}[N_T - 1] \end{bmatrix} \quad (12)$$

which aids in the construction of a model for the continuous stream with embedded transients. Using this sequential form, our model for the interference component of equation (3) becomes:

$$\hat{\mathbf{y}}[n] = \sum_{i=1}^{N_S} \sum_k \mathbf{U}_i[n-k] \mathbf{a}_i[k] \quad (13)$$

For values of the indices of $\mathbf{U}_i[\cdot]$ outside of the interval $[0, \dots, N_T - 1]$, we consider $\mathbf{U}_i = \mathbf{0}$. Our objective in cancelling the interference is to estimate the weight sequences $\mathbf{a}_i[k]$ so as to minimize the energy in the residual sequence:

$$\sum_{n \in I} (\mathbf{x}[n] - \hat{\mathbf{y}}[n])^2 \quad (14)$$

over some suitably chosen interval I .

2.1.3 Algorithms

Our algorithms operate on finite time intervals. In this section, we define as \mathbf{x}_n a segment of continuous data consisting of $N > N_T$ time steps and beginning at time n :

$$\mathbf{x}_n = \begin{bmatrix} \mathbf{x}[n] \\ \mathbf{x}[n+1] \\ \vdots \\ \mathbf{x}[n+N-1] \end{bmatrix} \quad (15)$$

The convolution operation defining our model of interference (from one source) then can be described by a matrix product:

$$\hat{\mathbf{y}}_n = \begin{bmatrix} \mathbf{U}[0] & \mathbf{0} & \mathbf{0} & \cdots & \mathbf{0} \\ \vdots & \mathbf{U}[0] & \mathbf{0} & \cdots & \vdots \\ \mathbf{U}[N_T-1] & \vdots & \mathbf{U}[0] & \cdots & \mathbf{0} \\ \mathbf{0} & \mathbf{U}[N_T-1] & \vdots & \cdots & \mathbf{0} \\ \mathbf{0} & \mathbf{0} & \mathbf{U}[N_T-1] & \cdots & \mathbf{U}[0] \\ \vdots & \vdots & \vdots & \vdots & \vdots \\ \mathbf{0} & \mathbf{0} & \mathbf{0} & \mathbf{0} & \mathbf{U}[N_T-1] \end{bmatrix} \begin{bmatrix} \mathbf{a}[n] \\ \mathbf{a}[n+1] \\ \vdots \\ \mathbf{a}[n+N-N_T] \end{bmatrix} \quad (16)$$

which simplifies the description of our first algorithm.

For a succinct description, we define the weight vector:

$$\mathbf{a}_n = \begin{bmatrix} \mathbf{a}[n] \\ \mathbf{a}[n+1] \\ \vdots \\ \mathbf{a}[n+N-N_T] \end{bmatrix} \quad (17)$$

of dimension $(N - N_T + 1) \cdot d \times 1$ and the model matrix:

$$\mathbf{M} = \begin{bmatrix} \mathbf{U}[0] & \mathbf{0} & \mathbf{0} & \cdots & \mathbf{0} \\ \vdots & \mathbf{U}[0] & \mathbf{0} & \cdots & \vdots \\ \mathbf{U}[N_T-1] & \vdots & \mathbf{U}[0] & \cdots & \mathbf{0} \\ \mathbf{0} & \mathbf{U}[N_T-1] & \vdots & \cdots & \mathbf{0} \\ \mathbf{0} & \mathbf{0} & \mathbf{U}[N_T-1] & \cdots & \mathbf{U}[0] \\ \vdots & \vdots & \vdots & \vdots & \vdots \\ \mathbf{0} & \mathbf{0} & \mathbf{0} & \mathbf{0} & \mathbf{U}[N_T-1] \end{bmatrix} \quad (18)$$

of dimension $N_c \cdot N \times (N - N_T + 1) \cdot d$. With these definitions our estimate of the interference is:

$$\hat{\mathbf{y}}_n = \mathbf{M} \mathbf{a}_n \quad (19)$$

Under the simplifying assumption that \mathbf{s}_n and $\boldsymbol{\eta}_n$ can be modeled as a uniform, independent, white Gaussian processes, a solution to the estimation problem can be found by choosing \mathbf{a}_n to minimize the squared residual norm:

$$\|\mathbf{x}_n - \mathbf{M} \mathbf{a}_n\|_2^2 \quad (20)$$

Continuous-source algorithm

Our first algorithm minimizes (20) directly, with the solution:

$$\mathbf{a}_n = (\mathbf{M}^T \mathbf{M})^{-1} \mathbf{M}^T \mathbf{x}_n \quad (21)$$

This least-squares problem minimization is well-posed, provided the matrix:

$$\mathbf{R} = \mathbf{M}^T \mathbf{M} \quad (22)$$

is full-rank. Note that \mathbf{R} is banded, and, in fact, block Toeplitz:

$$\mathbf{R} = \begin{bmatrix} \mathbf{R}_0 & \mathbf{R}_1 & \cdots & \mathbf{R}_{N_T-1} & \mathbf{0} & \cdots & \mathbf{0} \\ \mathbf{R}_{-1} & \mathbf{R}_0 & \mathbf{R}_1 & \vdots & \mathbf{R}_{N_T-1} & \ddots & \vdots \\ \vdots & \mathbf{R}_{-1} & \ddots & \mathbf{R}_1 & \vdots & \ddots & \mathbf{0} \\ \mathbf{R}_{-N_T+1} & \cdots & \mathbf{R}_{-1} & \mathbf{R}_0 & \mathbf{R}_1 & \ddots & \mathbf{R}_{N_T-1} \\ \mathbf{0} & \mathbf{R}_{-N_T+1} & \cdots & \mathbf{R}_{-1} & \mathbf{R}_0 & \ddots & \vdots \\ \vdots & \ddots & \ddots & \ddots & \ddots & \mathbf{R}_0 & \mathbf{R}_1 \\ \mathbf{0} & \cdots & \mathbf{0} & \mathbf{R}_{-N_T+1} & \cdots & \mathbf{R}_{-1} & \mathbf{R}_0 \end{bmatrix} \quad (23)$$

Where, for $0 \leq p \leq N_T - 1$:

$$\mathbf{R}_p = \sum_{j=p}^{N_T-1} \mathbf{U}^T[j] \mathbf{U}[j-p] \quad (24)$$

Note that the \mathbf{R}_p are dimensioned $d \times d$ and, by construction, $\mathbf{R}_0 = \mathbf{I}_{d \times d}$. Also, \mathbf{R} is symmetric, with:

$$\mathbf{R}_{-p} = \mathbf{R}_p^T \quad (25)$$

The core calculation is the evaluation of equation (21), which we break into two parts. The first is evaluation of $\mathbf{C} = \mathbf{M}^T \mathbf{y}$, which, by inspection, is a matrix convolution operation, and can be accomplished efficiently with overlap-add algorithms [Oppenheim and Shafer (1975)]. Given the result of this operation, the second part is evaluation of $\mathbf{a} = \mathbf{R}^{-1} \mathbf{C}$. Since \mathbf{R} is block Toeplitz, there is an efficient Levinson-Durbin recursion for the solution of this problem, which does not require construction of the full \mathbf{R} matrix or its inverse. For the reader who is not familiar with the Levinson-Durbin recursion, a derivation is provided for convenience in Appendix A.

A note about implementation: the matrices of (18) and (23) are enormous in any realistic problem and cannot be constructed explicitly. Consequently, the block Toeplitz structure must be exploited with the Levinson-Durbin recursion or some similar algorithm. Even exploiting that structure, the solution is computationally intensive on today's servers. For example, for a rank-3 subspace representation and a data block size of 3 minutes at 40 samples per second, a template size of 10 seconds, the system to be solved is

$(3 \cdot 170 \cdot 40 = 20,400) \ 20,400 \times 20,400$. For continuous stream processing we found it necessary to break the data into overlapping blocks, process each block separately and piece the results back together. We use three minute blocks, with one minute overlaps on each end, retaining the residual only for the central minute of each block. With this much overlap we do not notice any discontinuities in the processed residual segments when the continuous result is assembled. Although this method of processing contains a significant number of redundant

calculations, processing larger blocks to reduce redundancy (say 7 minute blocks with the same overlap) is not more efficient, because the cost of solution by Levinson-Durbin recursion is

$$O(N^2).$$

Another point regarding the continuous source algorithm is that when more than one source of interference is present, with model matrices \mathbf{U}_i , we simply concatenate the model matrices to form a combined model $\mathbf{U} = [\mathbf{U}_1 \ \mathbf{U}_2 \ \dots]$.

Intermittent-source algorithm

As noted in the introduction, a much faster solution is possible by approximating the excitation at each interfering source as a collection of discrete events. Instead of estimating the entire source time history as a continuous stream, we treat it as an intermittent, sparse sequence of excitations. We then perform a waveform fit and subtraction for distinct excitation events.

Referring to the data model of equation (3), we represent each of the continuous sources $\mathbf{f}_i[n]$ as a superposition of M_i impulsive forcing functions (corresponding to discrete interfering events) at distinct time indices $\{n_{ij}\}$:

$$\mathbf{f}_i[n] = \sum_{j=1}^{M_i} \mathbf{f}_{ij} \delta[n - n_{ij}] \quad (26)$$

In our empirical representation for the interference (equation 13) we treat each weight sequence as a similar superposition of delta functions:

$$\mathbf{a}_i[n] = \sum_{j=1}^{M_i} \mathbf{a}_{ij} \delta[n - n_{ij}] \quad (27)$$

The consequent interference model due to one source (i) is:

$$\sum_k \mathbf{u}_i[n-k] \sum_{j=1}^{M_i} \mathbf{a}_{ij} \delta[k - n_{ij}] = \sum_{j=1}^{M_i} \mathbf{u}_i[n - n_{ij}] \mathbf{a}_{ij} \quad (28)$$

The interference over all N_S sources is:

$$\hat{\mathbf{y}}[n] = \sum_{i=1}^{N_S} \sum_{j=1}^{M_i} \mathbf{u}_i[n - n_{ij}] \mathbf{a}_{ij} \quad (29)$$

Under this formulation, our cancellation objective is to find the collections of indices $\{n_{ij}\}$ and weights $\{\mathbf{a}_{ij}\}$ that minimize the functional:

$$\sum_n \left(\mathbf{x}[n] - \sum_{i=1}^{N_S} \sum_{j=1}^{M_i} \mathbf{u}_i[n - n_{ij}] \mathbf{a}_{ij} \right)^2 \quad (30)$$

We solve this problem with a type of alternation approach. We first estimate the indices, then the weights, then iterate the process. To estimate the indices, we perform a detection step, running a subspace detector [Harris, 2006; Harris and Dodge, 2011] over the continuous multichannel stream for each distinct interference source. Subspace detectors generate continuous detection statistics that range between 0 and 1, much like correlation coefficients.

Each detector operates by running a sliding window of length N_T over the continuous stream sample-by-sample, projecting the data in each successive window position into the subspace defined by the interference model representation \mathbf{U}_i . The detection statistic for a given time step is the energy of the projection normalized by the energy of the data in the window. The maximum possible projected energy is the energy in the data itself, so the maximum value of the ratio is one and equals one only when the data fall completely within the subspace (i.e. the trace in the window is a target interference waveform). We note, that if the model representation has rank one (i.e., if the representation consists of a single waveform), then the subspace detection statistic is exactly the square of the sample correlation coefficient between the data and the interference waveform.

The indices are estimated as the points at which the detection statistics exceed some threshold value. When several interference sources are in play, the detection process may be complicated by near simultaneous detections of an interference event by two or more detectors. The solution to this problem is to collect near-simultaneous triggers from all detectors and promote the one with the maximum statistic as the single detection. This reconciliation process determines the value of i in n_{ij} .

Once the indices of interference events are determined, the weights $\{\mathbf{a}_{ij}\}$ are estimated by minimizing the functional of equation (30). This process is simple if interfering events are sparse and do not overlap. In this case, interference events can be treated independently; the weights are obtained by a simple projection of the waveform data in the trigger window onto the subspace:

$$\mathbf{a}_{ij} = \begin{bmatrix} \mathbf{U}_i[0] \\ \mathbf{U}_i[1] \\ \vdots \\ \mathbf{U}_i[N_T - 1] \end{bmatrix}^T \begin{bmatrix} \mathbf{x}[n - n_{ij}] \\ \mathbf{x}[n - n_{ij} + 1] \\ \vdots \\ \mathbf{x}[n - n_{ij} + N_T - 1] \end{bmatrix} \quad (31)$$

For very active interference sources, such as the ones shown in Figure 4, interference waveforms frequently are superimposed. Then the estimation procedure is complicated by the need to estimate several weight vectors simultaneously over the union of time intervals of the interference events. In addition, with overlapped events, the subspace detection statistics may fall below the detection threshold for at least one of the events. Interference transients, even those well-characterized by one of the interference models, can be missed in the detection step. It is this effect which forces an iterative approach.

Our complete algorithm consists of estimating the indices $\{n_{ij}\}$ and weights $\{\mathbf{a}_{ij}\}$ in each iteration incrementally from the cancellation residual of the previous iteration. It often happens that cancellation of an interference transient in one pass of the algorithm exposes a new interference waveform to be detected in the next pass. In each pass, we add the indices of new detections to the accumulating set of indices for all events and re-estimate the weights for all detections. This process continues until no new detections are obtained. In the examples we consider, as many as four iterations produce usable detections.

Post-Processing Correction to the Continuous-Source Canceller

In the above, we have described two procedures for obtaining an estimate of the interference, $\mathbf{y}[n]$, resulting from a set of repeating transient noise sources: a continuous source algorithm and an intermittent source algorithm. Figure 6 displays a schematic picture of the behavior typically observed by the two algorithms in our case study of the SPITS icequakes. In the top trace (a) we see an original waveform containing a signal of interest in addition to background noise and high amplitude transients that we would like to remove from the data stream. In trace (b) we see the model for $\mathbf{y}[n]$ generated by the intermittent source estimator. This estimate is zero for all time samples except for those immediately following subspace detections; it is non-zero wherever there are sufficiently good matches with the interference source templates. Only two of the three unwanted transient signals are modelled in this example. While the subspace detector allows some variability in the source-time function, the sequence of multiple events which led to the superposition seen here results in a detection statistic lower than our predetermined threshold. As a consequence, the composite signal is ignored. Trace (c) is the residual, $\mathbf{s}[n] + \boldsymbol{\eta}[n]$, obtained from our observation $\mathbf{x}[n]$ by subtracting the interference estimate, i.e.:

$$\mathbf{s}[n] + \boldsymbol{\eta}[n] = \mathbf{x}[n] - \mathbf{y}[n] \quad (32)$$

The two transients present in (b) are cancelled with the rest of the waveform unchanged. As noted in the previous section, sometimes composite signals can be partially cancelled in one pass of the intermittent source algorithm, with further cancellation in subsequent iterations.

Trace (d) shows the estimate of $\mathbf{y}[n]$ from the continuous source model. The algorithm assumes that the interference source is always present. As a consequence, this estimate generally is non-zero. All three of the interference transients are well represented in the model, even the one resulting from the composite event. However, the amplitude of $\mathbf{y}[n]$ is a significant fraction of the amplitude of the desired signal $\mathbf{s}[n]$, such that when the subtraction is performed (trace e) that signal has been corrupted. Observation of long durations of data indicate that the amplitude of $\mathbf{y}[n]$ is a relatively constant fraction of the amplitude of $\mathbf{x}[n]$ for a given interference template (with characteristics controlled by, for example, time-bandwidth product of the interference waveform, number of channels observing). A greater number of channels (e.g. with a 3-component array rather than a vertical-only array) or a more complex transient

signal will match windows of unrelated data less well, reducing this fraction. Distortion of the original waveform is particularly prominent when a sharp contrast in amplitudes occurs over a time-window that is short in comparison with the length of the interference template. This is particularly problematic since it is prone to occur exactly at the onset of high SNR seismic phase arrivals.

Fortunately, it is relatively easy to identify the data segments for which it is beneficial to subtract $\mathbf{y}[n]$ from $\mathbf{x}[n]$. If the transient model $\mathbf{y}[n]$ represents the observation $\mathbf{x}[n]$ well at a given time, the match is good both in form and in amplitude. As a post-processing correction to the continuous source cancellation algorithm, we seek a scalar windowing function, $\mathbf{W}[n]$, such that

$$\mathbf{s}[n] + \boldsymbol{\eta}[n] = \mathbf{x}[n] - \mathbf{W}[n]\mathbf{y}[n] \quad (33)$$

with $\mathbf{W}[n] = 1$ where cancellation should be performed and $\mathbf{W}[n] = 0$ where the original

waveform should be preserved. We begin by setting $\mathbf{W}[n] = 1$ for all samples n and we

evaluate the instantaneous envelope functions $|\mathbf{x}[n]|$ and $|\mathbf{y}[n]|$ using the data traces and the

corresponding Hilbert transforms. For any sample n where either ratio $|\mathbf{x}[n]|:|\mathbf{y}[n]|$ or

$|\mathbf{y}[n]|:|\mathbf{x}[n]|$ falls below a nominal threshold (e.g. 0.9) we set $\mathbf{W}[n] = 0$ which results in a binary trace which is predominantly unity over data windows containing a nuisance signal and predominantly zero elsewhere. There are small clusters of exceptions due to waveform variability and coincidental similarity over very short windows. Applying a median filter with a length of a few seconds eliminates these, resulting in a rectangular function that is unity over durations comparable to those of the interference transients and zero elsewhere. Adding a short-duration cosine taper to each of the zero-to-unity and unity-to-zero jumps in this function

provides a windowing function $\mathbf{W}[n]$ that does not result in discontinuities (Figure 6, trace f).

The lower two traces in Figure 6 display the windowed continuous source model and the resulting residual, in the interference transients are removed with no distortion of the signal at other times.

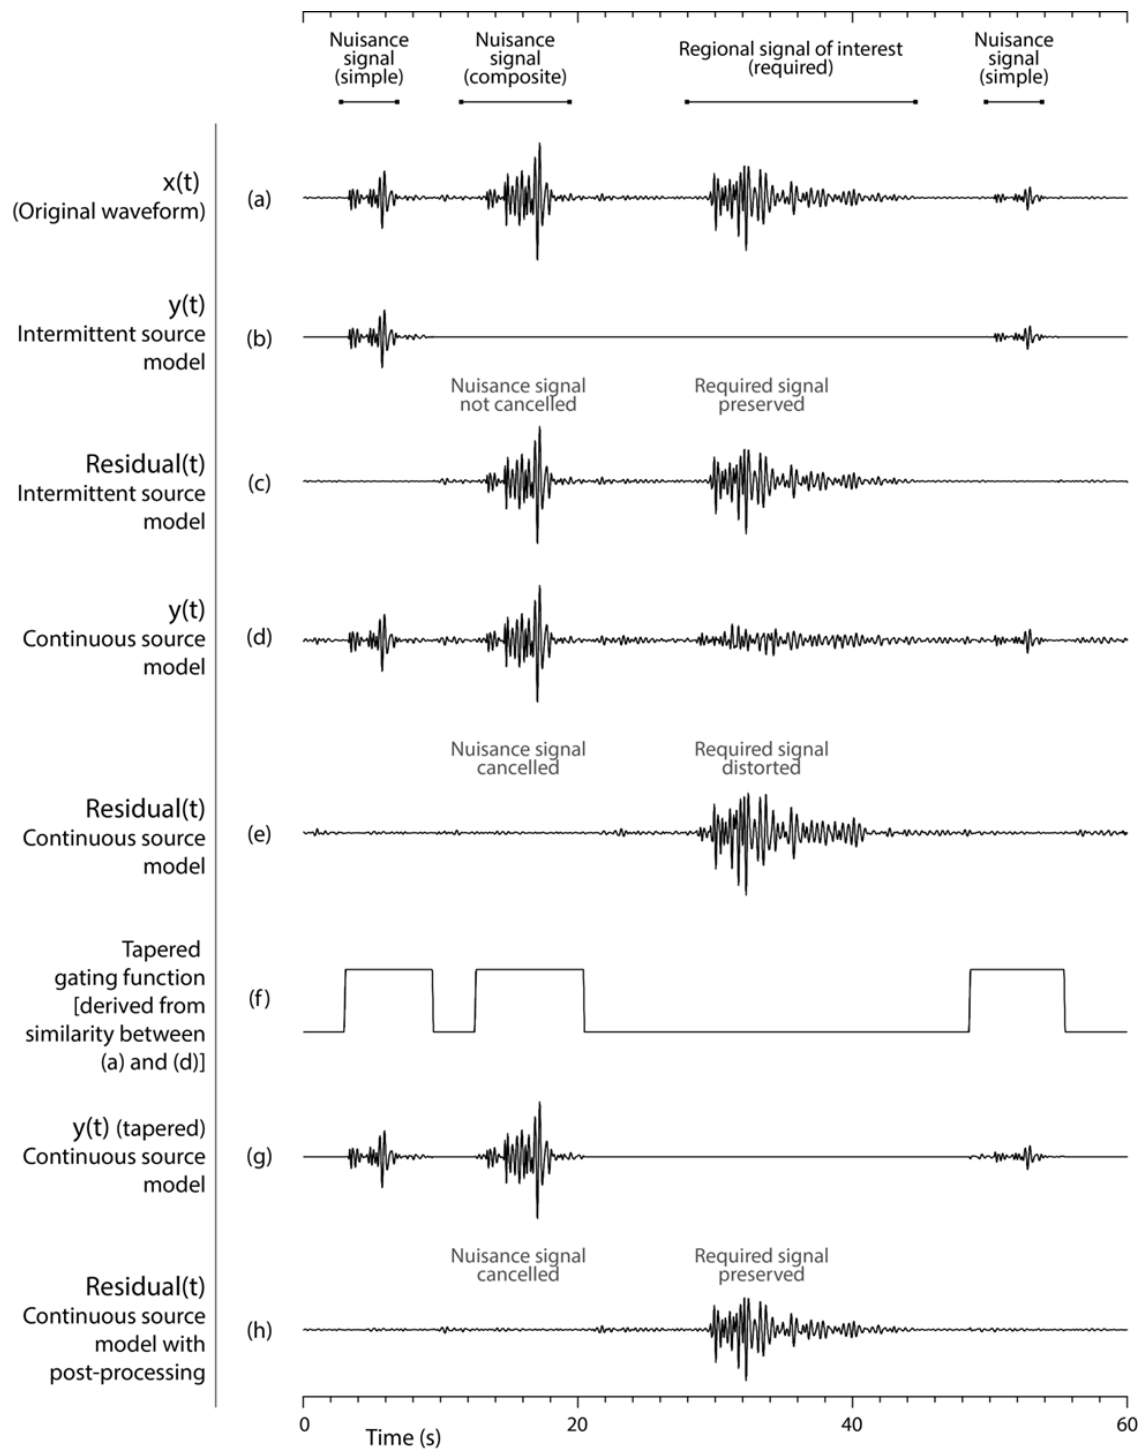


Figure 6 .Schematic depiction of transient signal cancellation using intermittent and continuous source models of the interference.

2.1.4 Spitsbergen Array Example

As an example of the application of the algorithms described in the previous section, we apply both to the continuous stream from the SPITS array during a period of prolific icequake occurrence. As described in the introduction, icequakes in the region surrounding the SPITS array dominate observations during periods of melting and freezing. P, S and Rg phases from these very local events are observed at the SPITS array and are, in many cases, reported by the automatic detection processor. Low propagation velocities, often in the range 1.3-2.0 km/s, are quite diagnostic for the Rg phases. However, for such nuisance P and S phases, the apparent velocities observed are indistinguishable from those estimated for P and S arrivals from regional events of monitoring interest. These nuisance detections cannot be eliminated from the detection stream without applying rigorous diagnostics to other properties of the signals (e.g. duration). Nuisance phases can constitute a very large fraction of all SPITS detections; failure to remove them from the processing pipeline can lead to spurious event hypotheses that slow the generation of seismic bulletins. For the year 2011 about 25% of SPITS detections appear to be slowly-propagating Rg phases. Figure 7 shows the azimuthal distribution of the 195,000 Rg-type phases reported in 2011, plotted at the geographical location of the SPITS array. We observe that the peaks of the azimuthal distribution coincide with the main directions towards neighboring glaciers.

During periods with high local event activity we often observe that “interesting” signals from regional or teleseismic distances are missed by the automatic detector because of data contamination by local signals. Even if nuisance transients do not overlap signals of interest, they may interfere with detection statistics, for example by reducing the Signal-to-Noise Ratio (SNR) for STA/LTA type detectors.




Figure 7. The rose diagram (red) shows the azimuthal distribution of 195,000 Rg-type phases reported in the SPITS detection lists for 2011. The rose diagram is centered at the location of the SPITS array.

As an example of cancellation performance, we apply both algorithms to continuous data from day 152 (June 1) of 2011. The SPITS array consists of 9 elements, a central element surrounded by a ring of 3 elements, with both surrounded by a larger concentric ring of 5 elements. The central element and the elements of the outer ring are three-component sensors. Elements of the inner ring are vertical sensors only. The data at this array are sampled at a rate of 80 samples per second. We apply the cancellers to two array configurations: the first consists of just the 9 vertical components, and the second consists of all 21 components.

The 5-step prescription described in section II was followed to identify the number of interference sources and to construct subspace representations for the waveforms they produce. In order to make a consistent comparison between the vertical and three-component data sets, the clustering and energy capture thresholds were manipulated in the design process to produce two clusters in each case, with a dimension of 3 for the subspace representing waveforms from the first cluster and dimension of 2 for the second subspace.

Results for the 21-channel cancellation example are shown in Figure 8 in a helicorder-style plotting format. At left in the figure is the original record from channel SPA1 BHZ for day 152 filtered into the 2 to 8 Hz band. It is clear from this panel that the amplitudes in the filtered waveform data are dominated by very large interference transients, such that events of interest

are difficult to identify visually. The middle panel of the figure shows the cancellation residual obtained with the continuous algorithm and appears to show a very substantial reduction in the total energy of interference events in the recording. Regional events of interest, unrelated to the interference transients now are the dominant signals in the display. The data were decimated to 40 samples per second prior to applying the continuous canceller. The continuous canceller is quite slow, completing cancellation of 24 hours of data in perhaps 6 hours on a 12-core Xeon server. This slow speed occurs despite the fact that the Java implementation is concurrent, making use of all 12 cores of the server.

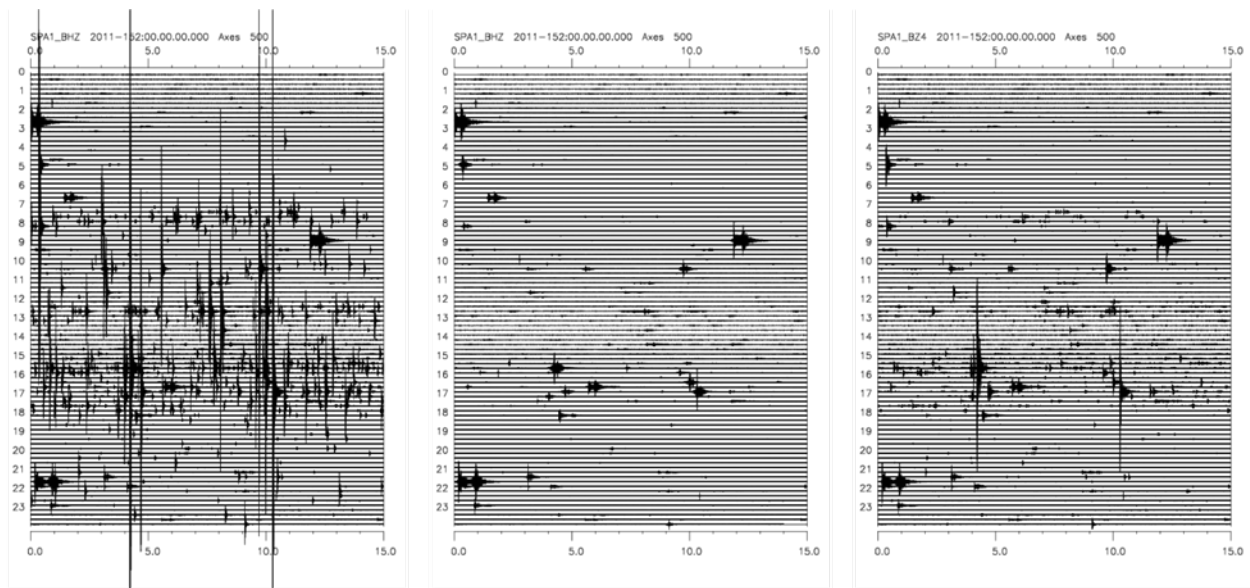


Figure 8. Heliograph type record of one channel (SPA1 BHZ) of the SPITS array (left), and cancellation residuals for the continuous canceller algorithm (middle) and the detecting canceller algorithm (right). Icequake removal is more thorough with the continuous algorithm, but at the price of greater computational cost.

The panel at right displays the residual obtained with the intermittent-source canceller after four iterations. The detection threshold was set to 0.6 in this case. Cancellation of the interference transients is not quite as thorough as for the continuous canceller, particularly for two rather large events and numerous small events. The small interference events probably were not detected, due to low SNR. The two large interference events may have been superimposed on regional waveforms. While cancellation may not have been quite as complete with this algorithm, it nonetheless has two distinct advantages: it is orders of magnitude faster than the continuous canceller and it has significantly less distortion of the remaining waveforms from regional events of interest. However, this latter advantage is eliminated by the post-processing correction to the continuous-source canceller.

For the intermittent-source canceller, it is illustrative to examine what is actually removed from the data stream in each iteration (Figure 9, bottom panels). It is clear that each iteration removes progressively fewer transients, and with lower amplitudes, than the previous iteration. However, since the computational cost of executing a single iteration of the detector is insignificant compared to the cost of the complete process, it is worth iterating the procedure until no subsequent change occurs in the waveform. In this example, this termination condition occurred after 4 iterations – the 5th iteration produced no further detections or cancellations.

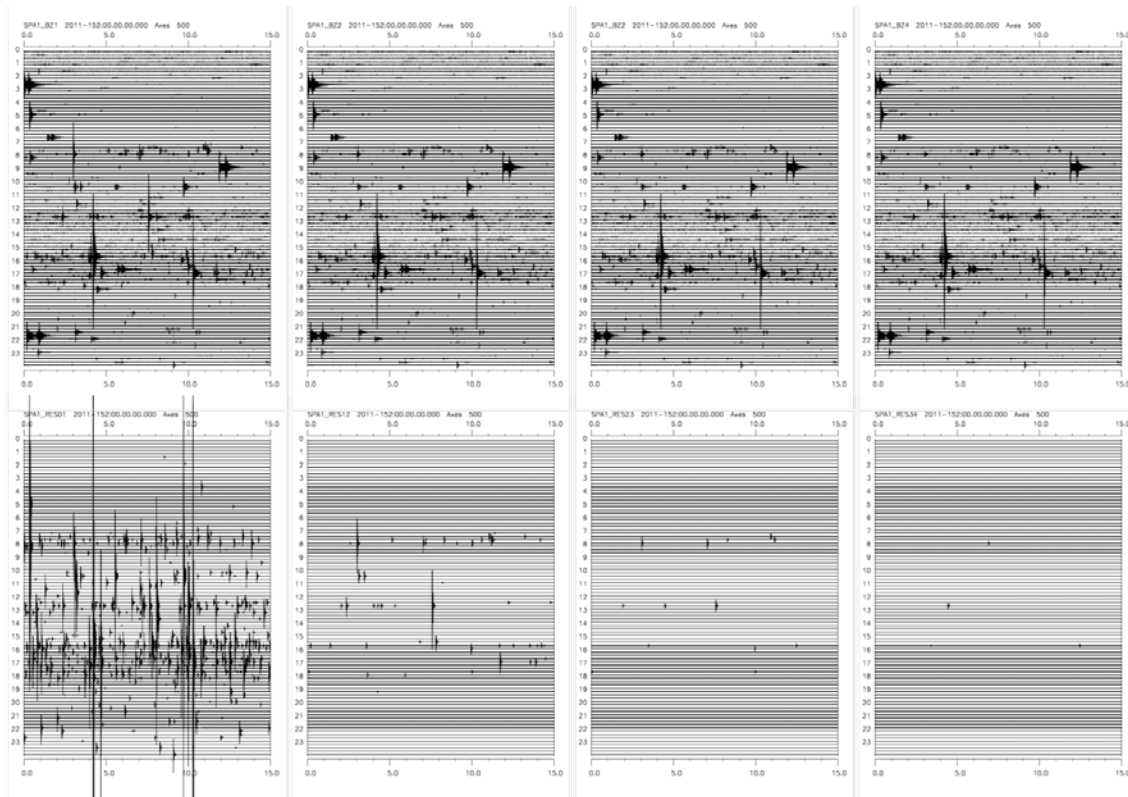


Figure 9. Composite of cancellation residuals for four iterations of the detection algorithm (top) and differences between the residuals from successive iterations (bottom). As anticipated, cancellation is subject to diminishing returns with iteration. The first two iterations produce significant numbers of detections and cancellations. There are no further cancellations after four iterations.

The advantages of the continuous source canceller are apparent in Figure 10, which compares details of the intermittent and continuous cancellation residuals over a 3-minute data segment. There are two clear icequake events in this time window. The intermittent source canceller reduces the amplitude of the first signal by an order of magnitude but fails to eliminate it to the point where it would not result in a detection subsequently with an STA/LTA detector. The second of the icequake signals is simply missed by the detector in the intermittent-source

algorithm. By contrast, both of the signals are reduced to the level of the background noise by the continuous-source canceller. The continuous-source cancellation residual has been modified by the post-processor designed to preserve desired signals. With the tapered window function, no changes are made to the waveform outside of the time intervals covering the icequake signals.




Figure 10. A 3-minute long data segment from the SPA0_BHZ channel. The top trace shows the full dynamic range of the filtered data and the lower 3 traces are all displayed to the same vertical scale. Traces 2, 3, and 4 from the top show respectively the original filtered data at higher gain, the cancellation residual from the intermittent-source algorithm, and the residual from the continuous-source algorithm.

The next 3 figures provide additional information on the necessity of the post-processing correction to the continuous-source canceller. These figures show a detailed section (1200 seconds) of the stream and various cancellation residuals. Figure 11 displays the results of cancellation using the 9 vertical elements of the SPITS array. The design process for the

cancellation algorithms resulted in two clusters of interference waveforms; one was represented with a subspace template of dimension 2 and the other of dimension 3. The top trace of the figure displays the original data from channel SPA0 BHZ. It contains a number of icequakes and a large regional event waveform.




Figure 11. Cancellation options for the 9-element array of SPITS vertical sensors show that it is possible to suppress icequakes without distorting desired regional signals. The top trace is an original filtered signal. The second and third traces are cancellation residuals for the intermittent-source and continuous-source algorithms respectively. The fourth trace is the gating function generated by the post-processor, and the last trace is the result of correcting the continuous-source residual with the gating function.

The second trace is the residual from the intermittent-source canceller (also called detection canceller) using the two interference source representations of dimensions 2 and 3. It largely eliminates or suppresses the icequakes without distorting the regional waveform. The third trace is the residual from the continuous-source canceller, which is better at suppressing the icequakes, but at the cost of distorting the regional waveform. At this scale, the distortion is most obvious in the loss of amplitude of the P and S phases. The fourth trace is the gating function generated by the post-processor indicating where the cancellation residual should be stitched into the original stream. The final trace is the corrected continuous-source residual,

showing that the gating function preserves the regional waveform while allowing cancellation of the icequakes. This residual is the cleanest of the cancellation options.




Figure 12. Cancellation residuals for the full array of 21 vertical and horizontal elements of the SPITS array show that larger numbers of channels mitigate distortion. The traces are the equivalents of those in Figure 11 for cancellation performed on 21 channels instead of 9.

Figure 12 provides cancellation residuals from the full 21-element SPITS array for comparison with Figure 11. In this case the cancellers also used two sources with 2- and 3-dimension subspace representations. Since the total number of degrees of freedom in the cancellers is that same as the prior example, but the number of channels is greater, it may be anticipated that constraints on distortion of the regional signal are improved. This is the case, as can be seen in the third trace of the Figure 12. Amplitude loss in the regional event is less pronounced in this figure. It also is apparent that the icequakes are suppressed less effectively by the intermittent-source algorithm (trace 2) when 21 channels are used instead of 9. This effect is a consequence of the larger number of constraints that have to be met in fitting 21 channels of data with 5 degrees of freedom instead of 9 channels. As before, the continuous-source cancellation residual is the best among the three cancellation options.

Detail of the distortion created by the continuous canceller and corrected by the gating function of the post-processor is displayed in Figure 13. The top trace is again the original filtered data from station SPA0 BHZ. The next two traces are the residuals from the continuous-source

cancellers acting on 9 channels and 21 channels of data respectively. This distortion manifests itself in amplitude reduction in the waveforms, but more disturbingly in P-wave precursor artifacts. Distortion definitely is reduced with the addition of more channels to the processed data stream (comparing trace 3 to trace 2). However, the post-processing algorithm restores the regional signal (bottom trace) without adding noticeable artifacts to the rest of the stream. With this observation, the continuous-source canceller, corrected by the post processor appears would be the algorithm of choice were it not for the fact that the continuous canceller is several (~ 3) orders of magnitude slower than the intermittent-source algorithm.




Figure 13. Details of cancellation residuals for the three algorithms around the regional waveform shown in Figures 11 and 12. The top trace is the original SPA0 BHZ data, filtered. The next two traces are the residuals from the continuous-source cancellers using 9 and 21 channels of the SPITS array, respectively. The bottom trace is the residual from the continuous-source canceller corrected by the post-processor. Waveform distortion is avoided in this case.

2.1.5 Summary

In this chapter we have proposed a collection of algorithms for removing repetitive nuisance transients from array data streams. The algorithms are applicable whenever an array is surrounded by sources producing highly correlated, short-duration transient waveforms. At the

moment, we have identified one practical algorithm, the intermittent-source canceller, which can be applied to a fairly large number of sources surrounding an array and is sufficiently fast to be embedded in a monitoring pipeline. We have, in addition, examined a continuous-source canceller, which appears to be more effective in removing nuisance transients, but is more computationally demanding and requires a post-processing step to prevent it from distorting signals from events of interest. The intermittent-source canceller is asymptotically equivalent to the continuous-source algorithm in the sense that it tends to the latter algorithm as the threshold in the detection step is reduced to zero.

We will report separately (next chapter) on a study quantifying the performance of the intermittent-source canceller, but state here that it appears to be highly effective in reducing the number of interference detections that currently create bottlenecks in pipeline associators. Our expectation is that next-generation pipelines could include adaptive cancellers for each array that would clean the array data streams of locally-generated transient interference. We envision hierarchical adaptive pipelines that would contain, at their base levels, adaptive mini-pipelines designed to detect repetitive nuisance events at each array and generate, then apply, cancellers automatically. The research that remains to be done in order to achieve this vision is for the autonomous identification of repetitive transients in a continuously adapting pipeline preprocessor or lower-level component of the hierarchy. What we have done here is to demonstrate that cancellation is technically feasible given a set of templates that represent a majority of interfering transients.

2.2 A Revised Detection Framework Architecture

2.2.1 Introduction

We significantly revised the data handling architecture in adapting the detection framework [Harris and Dodge, 2011; Kværna et al., 2014] to perform signal cancellation, as described in the previous section. In part this was because the previous strategy of sharing data buffers among multiple objects made it more difficult to implement a process that would be modifying those data buffers. More importantly, as we move towards development of a detection framework that can scale out a cluster of computers, we can no longer take advantage of shared memory. In the new architecture, the various computational units no longer have any shared state. Instead, they pass data messages to one another.

Figure 14 shows the flow of data through the current architecture. Data is produced by the StreamServer which is responsible for assembling data into blocks of the required size. The StreamServer produces objects of type StreamSegment which contain the time series data, and the necessary metadata to allow correct routing. In the current architecture, the StreamSegment objects are passed through a BlockingQueue to an object called the StreamProcessor. Within the StreamProcessor 0-N StreamModifier objects may consume a StreamSegment and emit a different StreamSegment. Currently, there are two possible StreamModifiers (DownSampler and SegmentedCancellor). The final modified StreamSegment is passed to a StreamTransformer which emits a TransformedStreamSegment. This is a StreamSegment that contains in addition a DFT of itself.

At this point, a parallel section is entered. Within this section, for each Detector:

1. The TransformedStreamSegment is used to create a DetectionStatistic
2. The DetectionStatistic is consumed by a DetectionStatisticScanner which emits
3. A Collection of TriggerData.

At this point, the threads join and all the TriggerData is ingested by a TriggerProcessor. The TriggerProcessor archives the triggers and produces a collection of Trigger objects which contain additional state created by the TriggerProcessor.

The Triggers are then passed to the DetectionProcessor which applies the necessary rules to promote some triggers to Detections. Any such Detections are then handed off to a DetectorCreator and to a DetectionArchiver.

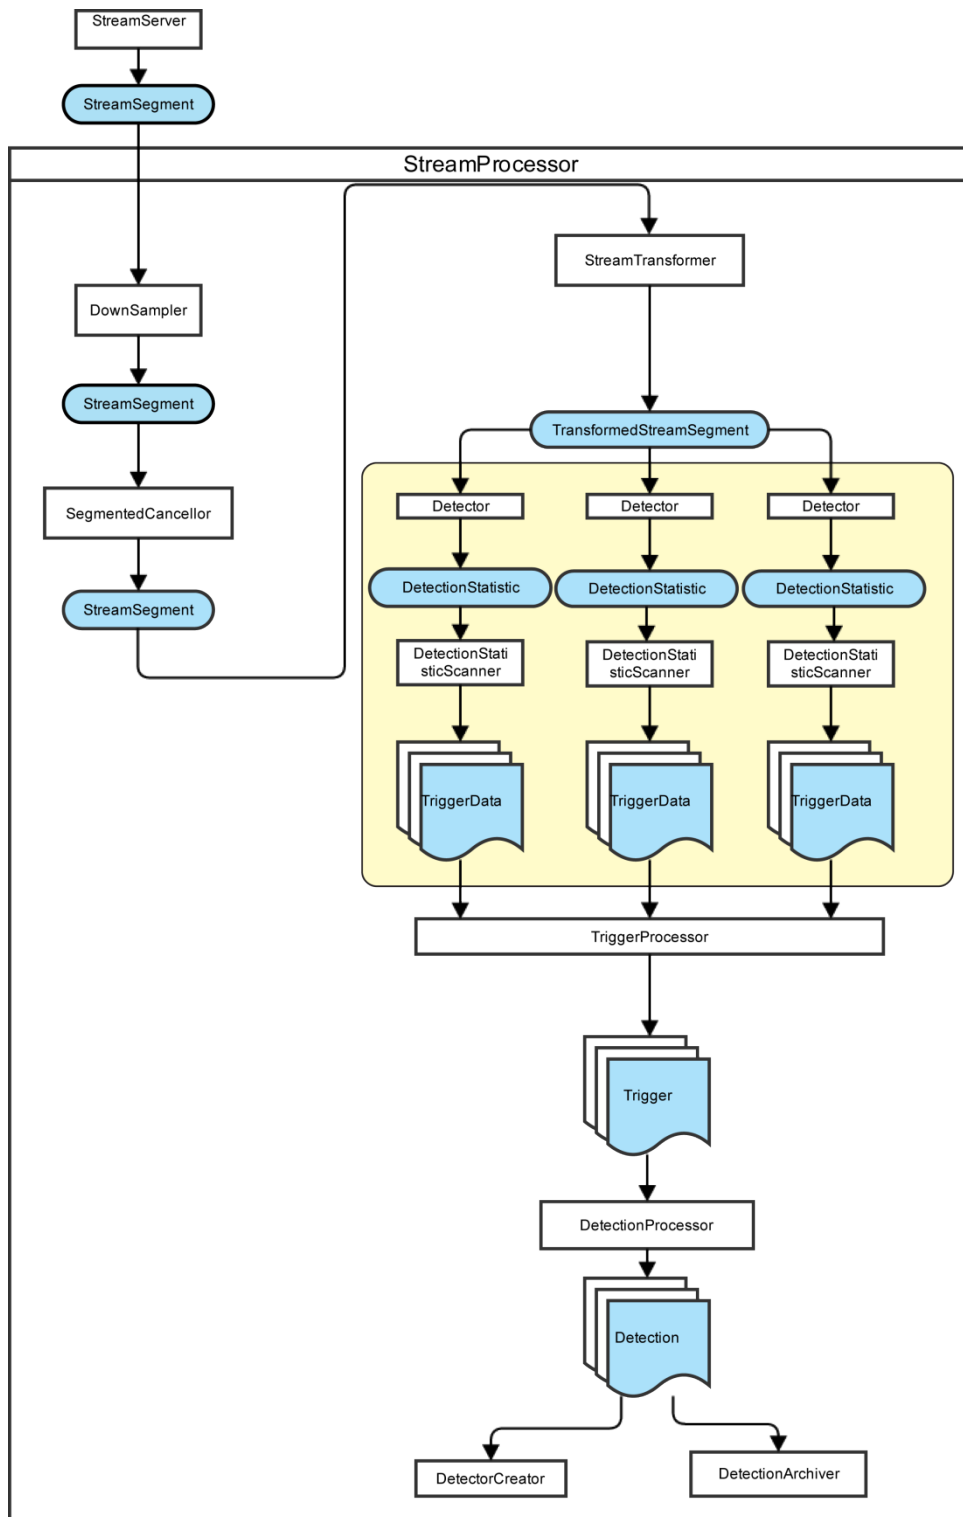


Figure 14. Flow of data through the revised detection framework architecture.

2.2.2 Test of the Signal Cancellation Component in the Detection Framework

To test the cancellation option, the framework was run against a 3-day segment (day 2011151 through day 2011153) of vertical-component data from the SPITS array. During this time interval intense icequake activity had occurred, with hundreds of such signals visible at SPITS. Figure 15 shows the SPA0-BHZ channel for the time period with an inset showing the form of a typical signal.




Figure 15. 3 days of SPA0-BHZ data which is dominated by signals from nearby icequakes.

For reference, we first ran the detection framework without cancellation enabled. The framework was configured to use all the vertical-component channels filtered into the band 2-8 Hz and with a single STA/LTA spawning detector. New templates created by the framework were constrained to be between 10 and 30 seconds in length, and the detection threshold of the correlation detectors was 0.6.

This run produced 49 correlation detectors and a total of 291 detections. Two of the correlators accounted for 242 of the detections, and all of those signals appear to be due to icequakes. These detections are shown in Figure 16.




Figure 16. (Top) 94 detections produced by detector 92532 and (bottom) 148 detections from detector 92541 produced during the first run of the framework.

The remaining detections are distributed among the other detectors, with most having only a single detection. These appear to be mostly local earthquakes although a few apparent icequake signals are present as well. Figure 17 shows all these detections.




Figure 17. The 49 detections produced by the remaining detectors. The blue-shaded signals also appear to be icequake detections, but, perhaps because of interfering signals, were assigned to unique detectors.

Cancellation run (RUNID 365)

The discrete segmented cancellation processor provided by Deschutes Signal processing uses a two-step algorithm. First the signals to be cancelled are detected using subspace detectors. At each detection point, a best-fit to the templates is constructed and then subtracted from the signal. This process is iterated until no more detections are found. The detectors used by the canceller are produced by the cancellation software given an input set of nuisance signals. This processing is controlled by a number of parameters, some of which are shown below.

Table 1. Parameters controlling detector creation and cancellation behavior as they were set for this experiment.

templateLength	6.0
maxOffset	0.5
clusteringThreshold	0.9
minNumEvents	10
energyCapture	0.9
detectionThreshold	0.4
peakHalfWidth	0.2
simultaneityThreshold	2.0
numberOfIterations	5

For the execution of RUNID 365 the framework was configured identically to that of its previous run, except that the cancellation flag was set, and a group of 272 icequake signals was provided to serve as input for the signal canceller. From these, the software created two cancellation detectors. The effect of the cancellation can be seen by visual inspection of the signals pre- and post-cancellation. This is shown in Figure 18 below.




Figure 18. 3 days of SPA0-BHZ data prior to cancellation (top) and after cancellation (bottom). The traces are plotted at the same scale.

Comparison of Results

In total, the canceller performed 2047 signal cancellation operations. However, because the cancellation process is iterative, the number of independent icequake signals is smaller. To the

nearest second, there were 952 distinct signals cancelled, and to the nearest 10-seconds there were 817 distinct cancellations.

In contrast to the first run of the framework, RUNID 365 produced only 71 detections. Of these 17 were detected by two detectors (92584 and 92594) and are icequake signals. Table 2 shows these detections grouped by detector and associated by time with detections from the first run. All 17 are associated with detections by (RUNID 364) detectors 92541 and 92532, both of which were formed on icequake signals. Visual inspection of the detection segments confirms the signal type.

Table 2. The 17 detections from RUNID 365 that match (nearly) in time with icequake detections from RUNID 364.

364 detectorid	364 detectionid	365 detectorid	365 detectionid	Time difference
92541	195546	92584	195804	-6.6625
92541	195537	92584	195802	-6.6615
92532	195649	92584	195822	5.775
92541	195611	92584	195819	-6.675
92541	195654	92584	195824	-6.6625
92532	195514	92584	195795	5.769445
92541	195550	92584	195806	-6.6615
92541	195567	92584	195811	-6.6625
92541	195763	92594	195842	-4.9
92541	195685	92594	195829	-4.9
92541	195661	92594	195825	-4.8865
92541	195598	92594	195815	-4.9125
92541	195606	92594	195818	-4.9
92541	195604	92594	195816	-4.9
92532	195564	92594	195810	7.526389
92541	195651	92594	195823	-4.9125
92541	195582	92594	195813	-4.9

The cyan-shaded group is by detector 92584 and the green-shaded group is by detector 92594. The first two columns are the detector number and detection-id from the first run. The third and fourth columns show the corresponding information from the second run, and the last column is the difference in detection times.

Given the success of the signal canceller in removing so many other nuisance signals it was curious that these detections still remained. Also, why are the detection time differences so large? Both of these groups of detections were formed by detectors whose template included

small icequake signals followed by a very much larger icequake signal near the end of the template. The residual from the canceller was still a detectable signal, but the detailed structure was sufficiently different to change the point at which the STA/LTA detector could exceed its threshold. Figure 19 shows an example from the template event for detector 92584. In this example, two icequake signals are present. In run 1, the detection is declared at the onset of the first signal. But, after cancellation, the first signal is almost entirely removed, so the detection point shifts by several seconds.




Figure 19. This figure shows how partial cancellation can cause changes in detection time for cases when multiple instances of the nuisance signal are present in close proximity.

Of the remaining 54 detections in RUNID 365, 39 are exact matches in time to non-icequake detections from RUNID 364. There are 11 detections from RUNID 364 that were not seen in RUNID 365. Another 8 of the detections in RUNID 365 are on partially-cancelled signals. There are also 8 detections in RUNID 365 that were not seen in RUNID 364. The complete comparison

of the detection results between runs 364 and 365 is presented in Table 3 below. The unshaded rows summarize the exact matches showing the detectorid and detectionid from each run as well as the detection time.

Table 3. Accounting of all detections in RUNID 364 and RUNID 365 excepting 242 known icequake detections from RUNID 364 detectors 92532 and 92541 and 17 icequake detections from RUNID 365 detectors 92584 and 92594.

92517	195486	92567	195777	2011/05/31 (151) 06:06:14.000
92518	195487	92568	195778	2011/05/31 (151) 07:26:22.000
92519	195488	92569	195779	2011/05/31 (151) 17:36:25.000
92520	195489	92570	195780	2011/05/31 (151) 17:46:47.000
92521	195490	92571	195781	2011/05/31 (151) 17:57:37.000
92522	195491	92572	195782	2011/05/31 (151) 18:25:17.000
92523	195492	92573	195783	2011/05/31 (151) 19:30:49.000
92524	195493	92574	195784	2011/05/31 (151) 21:14:12.000
92525	195494	92575	195785	2011/05/31 (151) 21:47:19.000
92526	195495	92576	195786	2011/06/01 (152) 01:13:47.000
92527	195496	92577	195787	2011/06/01 (152) 01:32:20.000
92528	195497	92578	195788	2011/06/01 (152) 02:11:27.000
92529	195498	92579	195789	2011/06/01 (152) 02:29:58.000
92530	195499	92580	195790	2011/06/01 (152) 04:30:49.000
92531	195500	92581	195791	2011/06/01 (152) 05:54:13.000
92529	195501	92579	195792	2011/06/01 (152) 06:31:23.000
92533	195505	CANCELLED		2011/06/01 (152) 07:34:35.000
Detection on partial cancellation		92582	195793	2011/06/01 (152) 07:38:15.000
92534	195508	CANCELLED		2011/06/01 (152) 07:47:02.000
Detection on partial cancellation		92583	195794	2011/06/01 (152)

Table 3. Continued

				07:50:42.000
92535	195522	92585	195796	2011/06/01 (152) 08:56:50.000
92536	195523	92586	195797	2011/06/01 (152) 09:15:11.000
92537	195524	92587	195798	2011/06/01 (152) 09:16:42.000
92538	195525	92588	195799	2011/06/01 (152) 09:25:31.000
92539	195527	92589	195800	2011/06/01 (152) 09:34:06.000
92540	195528	92590	195801	2011/06/01 (152) 10:00:19.000
Correlation detection 2 of 3 (Large icequake signal masked it in run 1.)		92585	195803	2011/06/01 (152) 11:07:52.000
92542	195549	92591	195805	2011/06/01 (152) 12:09:55.000
Detection on partial cancellation		92592	195807	2011/06/01 (152) 12:28:16.000
Detection on partial cancellation		92593	195809	2011/06/01 (152) 12:37:06.000
Correlation detection based on partially-cancelled signal		92592	195808	2011/06/01 (152) 12:37:52.000
Detection of signal masked by icequake signal?		92595	195812	2011/06/01 (152) 13:00:06.000
92543	195586	92596	195814	2011/06/01 (152) 14:23:09.000
92544	195587	CANCELLED		2011/06/01 (152) 14:26:39.000
92545	195594	CANCELLED		2011/06/01 (152) 15:01:28.000
92546	195596	CANCELLED		2011/06/01 (152) 15:04:19.000
92547	195599	CANCELLED		2011/06/01 (152) 15:16:56.000
Detection on partial cancellation		92597	195817	2011/06/01 (152) 15:28:38.000
Detection on partial cancellation		92598	195820	2011/06/01 (152) 15:33:58.000
92548	195624	CANCELLED		2011/06/01 (152) 15:44:04.000
Signal onset masked by icequake signal		92599	195821	2011/06/01 (152) 16:35:34.000
92549	195659	CANCELLED		2011/06/01 (152) 16:58:31.000
92550	195668	92600	195826	2011/06/01 (152) 17:22:20.000
92551	195670	This is a 2nd detection on the same signal that only occurred in run 1		2011/06/01 (152) 17:22:21.000
92552	195671	92601	195827	2011/06/01 (152)

Table 3. Continued

				17:26:45.000
Detection on partial cancellation		92602	195828	2011/06/01 (152) 17:35:34.000
Signal preceded by partially-cancelled icequake signal		92603	195830	2011/06/01 (152) 18:04:15.000
92553	195704	CANCELLED		2011/06/01 (152) 19:57:35.000
92554	195708	92604	195831	2011/06/01 (152) 21:08:40.000
Correlation detection 3 of 3 (Large icequake signal in coda caused correlation to fail in run 1)		92579	195832	2011/06/01 (152) 21:45:08.000
92555	195719	92605	195833	2011/06/01 (152) 23:35:32.000
Detection of signal masked by icequake signal?		92606	195834	2011/06/02 (153) 01:05:27.000
Signal preceded by partially-cancelled icequake signal		92607	195835	2011/06/02 (153) 01:17:51.000
92556	195740	CANCELLED		2011/06/02 (153) 07:44:20.000
92557	195746	92608	195836	2011/06/02 (153) 10:00:17.000
Detection of signal masked by icequake signal?		92609	195837	2011/06/02 (153) 12:35:06.000
92558	195752	92611	195839	2011/06/02 (153) 12:39:49.000
92559	195754	92612	195840	2011/06/02 (153) 13:08:28.000
92560	195759	92613	195841	2011/06/02 (153) 16:39:04.000
92561	195765	92614	195843	2011/06/02 (153) 20:30:03.000
92535	195769	92585	195844	2011/06/02 (153) 21:25:06.000
92562	195771	92615	195845	2011/06/02 (153) 22:18:27.000
92563	195772	92616	195846	2011/06/02 (153) 22:25:35.000
92564	195773	92617	195847	2011/06/02 (153) 22:44:32.000
92565	195776	In Last Buffer. Not seen for detection		2011/06/03 (154) 00:02:06.000

The yellow shaded rows present information for cases when a detection occurred in run 364 but there was no corresponding detection in run 365. The green-shaded rows show information for cases when there was a detection in run 365 but no corresponding detection in run 364 and which are probably valid non-icequake signals. The pink-shaded rows are for cases when a detection was produced on a partially-cancelled icequake signal. The cyan-shaded rows are new detections of questionable validity.

Missing Detections in RUNID 365

There were 11 detections produced in RUNID 364 that were not seen in RUNID 365 and that were not in the two groups of known icequake signals. These are shown in yellow in Table 3 and the (un-cancelled) signals are shown below in Figure 20. All but the last of these signals are visually similar to the icequake signal shown in Figure 15, but they appear to have higher noise or other changes in structure sufficient that they did not correlate well enough with the templates of detectors 92532 or 92541 to be included in their detections. However, the multi-rank subspace detectors used by the cancellation processor were able to detect these signals, so they were cancelled sufficiently to not be detected in run 365.

The last signal is a special case. The time stamp of this detection is 2011/06/03 (154) 00:02:06.000. This is 2 minutes past the latest data the framework was requested to process, so it is almost certainly the last buffer of data pushed into the canceller. The framework simply shut down before this buffer was retrieved from the canceller. Although this is a missed detection in the context of this experimental run, in a continuously-operating system there is no “last buffer”, so this situation could not arise in practice.




Figure 20. Signals that were detected in the first run of the framework, but which were not detected in the run with cancelling turned on. The numbers are respectively the detector number and detection-id from run1. In all but the last case, the signals appear to be from icequakes, and were removed by the canceller. The last signal was not detected because it occurred in the very last data buffer processed by the system.

Detections on Partial Signal Cancellations

There are 8 cases of unwanted detections in run 365. In these cases, a detection occurred in run 365 with no corresponding detection in run 364, and the detection was on a partially-cancelled icequake signal. These cases are shown in Table 3 as the pink-shaded rows. Figure 21 shows the signals involved. (Only 7 signals are shown. The 8th was a correlation detection on the template of detector 92592.) Figure 21 shows the original signals in light gray and the cancelled signals on which the detections were made in magenta. The residual signals are likely to be from icequake sources since they maintain the same phase structure as the un-cancelled signals and have very little energy outside the bounds of the un-cancelled signals.

What makes these detections a little problematic is that the cancellation process and mixing of signals can produce residuals that have some resemblance to local earthquake seismograms. Clearly, an operational system that included cancellation would have to make it very easy for analysts to be able to observe the un-cancelled signal in context to help sort out these (hopefully rare) occurrences.

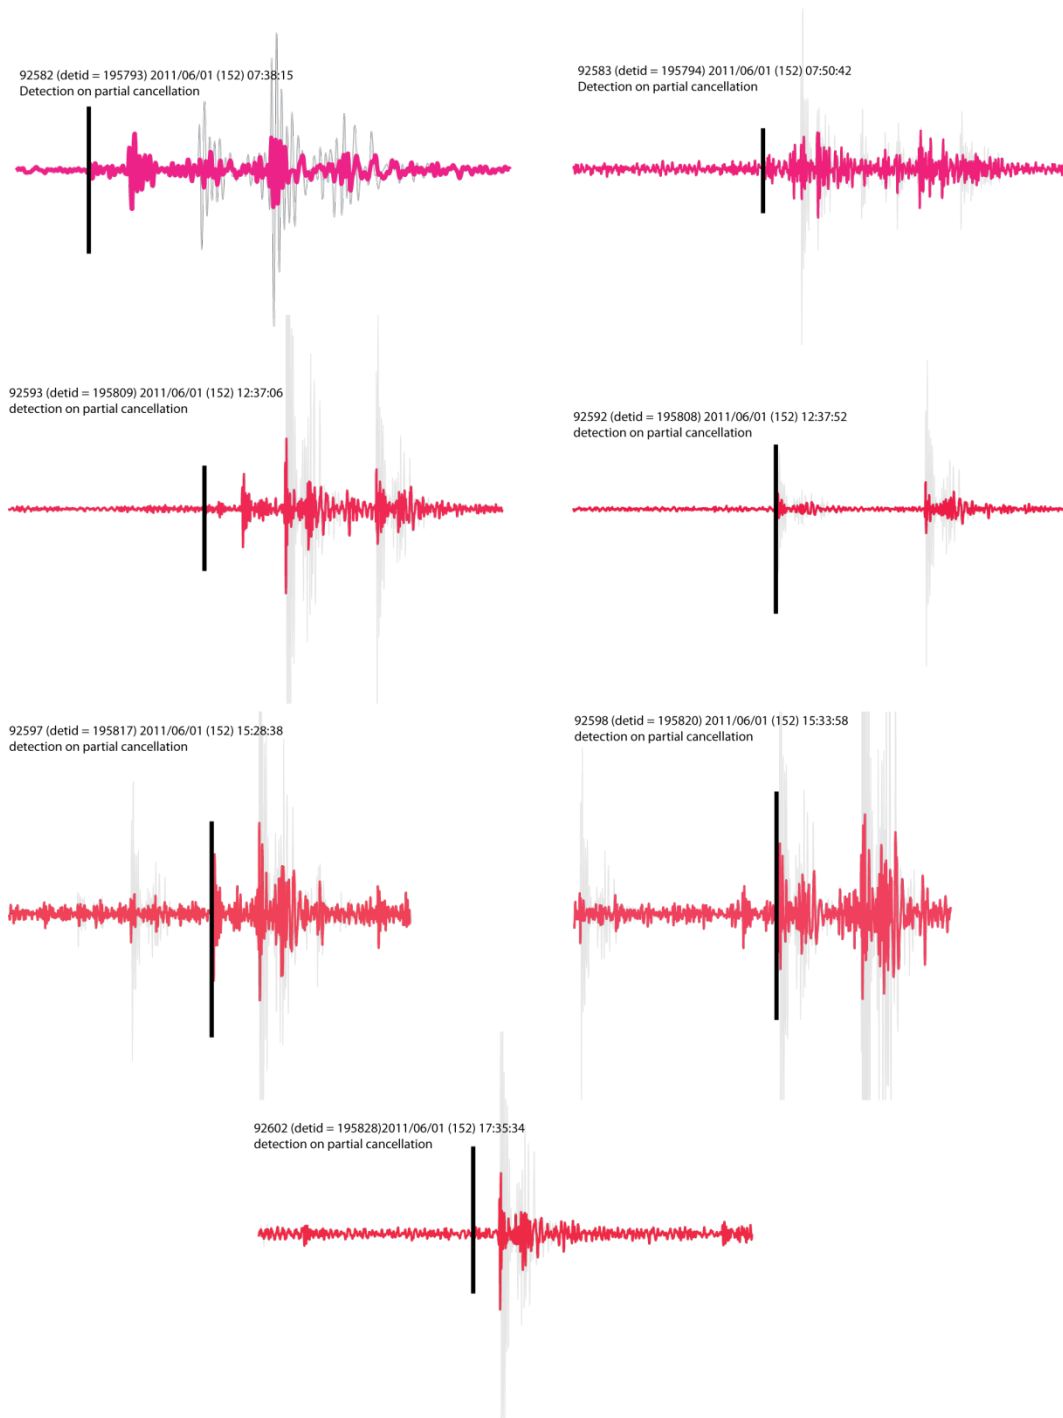


Figure 21. Instances in which the framework produced detections on incompletely cancelled icequake signals. The magenta traces are the post-cancellation signals, and the light gray shows the original signal. The detection time is indicated by the vertical black line.

New detections

There are also 5 cases where a detection was made in run 365 with no corresponding detection in run 364 and in which the detection is likely to be for a valid non-icequake signal. These are shown in Table 3 in green. In the cases shown in Figure 22 (detections 195821, 195830, 195835) a signal failed to be detected in run 364 because it fell within the blackout interval following an icequake signal detection.




Figure 22. The three cases where a valid signal was not detected in run 365 because the signal onset was in the blackout interval for an icequake detection.

After cancellation, in two cases the detection is still at the onset of the icequake signal, but the detection is by the STA/LTA detector instead of by one of the correlators. In the third case, the cancelled icequake signal trigger was discarded in favor of the trigger on the new signal. In two cases (detections 195803 and 195832) a potential correlation detection failed because of the presence of a much larger icequake signal. These are shown in Figure 23.




Figure 23. The two cases where a signal failed to be detected by correlation detectors because of the presence of a strong icequake signal.

In both cases, the signals shown in blue were detected by correlators in run 364, but the third signal was missed. In run 365, the cancelled icequake signal was sufficiently attenuated that the third signal was detected. Finally, there are three possible “new” detections (see Figure 24). The residual signals may be the remnants of the cancelled icequake signal, but they do have a shape very consistent with that of a local earthquake signal. Also, the first (detection 195812) begins shortly before the start of the cancelled icequake signal.

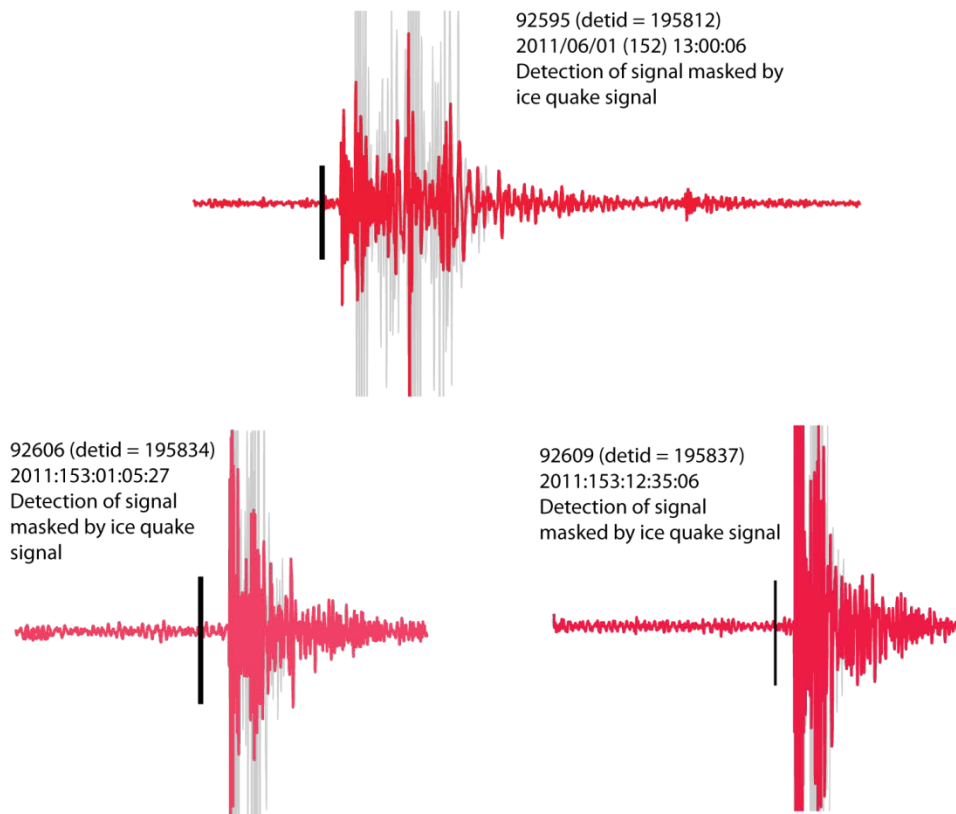


Figure 24. The three cases where a valid earthquake signal buried within an icequake signal may have been detected.

2.2.3 Summary

The intermittent-source cancellation algorithm has been implemented into the so-called Detection Framework. In order to facilitate the implementation of new methods and outscaling data processing to a cluster of computers, the Detection Framework was successfully upgraded with a new architecture.

The performance of the revised Detection Framework with respect to cancellation of local icequake signals was evaluated using data from the SPITS array. For a three day period the cancellation provided a factor 10 reduction in nuisance icequake signals. The data “cleaned” from nuisance signals again provide for improvements in both manual and automatic data analysis.

We conducted an experiment with automatic reprocessing of the “cleaned” SPITS data within the revised Detection framework. This demonstrated improvements in detection performance, but also revealed some issues to be resolved regarding partial-only signal cancellation of time-overlapping nuisance signals.

2.3 Empirical Matched Field Processing and Adaptive Beamforming

2.3.1 Introduction

The principal thrust of this project is to examine the potential for contextual signal processing to improve detection and association functions in processing pipelines. The idea is to adapt the signal processing algorithms – which today make unrealistically simple assumptions about the seismic background context in which they operate – to make more sophisticated use of present knowledge of context. In support of this idea, one of our objectives is to take a fresh look at adaptive beamforming algorithms for suppression of signals and noise that compete with signals of interest. Adaptive beamforming techniques use estimates of the structure (covariance) of background interference to suppress that interference. These techniques operate by reducing the response of arrays at wavenumbers where interference is strong, while simultaneously maintaining a fixed gain at wavenumbers from which desired (target) signals arrive. These techniques were developed for verification seismology in the 1970s, but were largely abandoned due to their sensitivity to signal model error. They assume that the spatial structures of target signals adhere closely to a plane wave model. Because the spatial structures of incident waves often are significantly non-planar (resulting from unmodeled refraction and scattering), adaptive techniques often suppress target signals as well as interference. In a sense, adaptive methods need to know what signal structure they are protecting as they aggressively configure an array's wavenumber response to suppress interference.

Recent developments in empirical detection methods, especially empirical matched field processing, suggest that signal models may be made sufficiently accurate to allow adaptive beamforming to function as intended. The approach, in keeping with modern trends of large scale waveform archival, is to use past events to provide target waveform calibrations. Empirical matched field processing estimates the spatial structure of target phases from prior observations in a large collection of narrow frequency bands. The quasi-frequency domain calibration that results minimizes sensitivity to variable or unknown source time histories.

As used here, matched field processing is closely related to frequency-domain beamforming, and, in fact, can be viewed as a calibrated form of beamforming. To describe the technique, we first examine a narrowband approximation to frequency-domain beamforming, then generalize the result. We start with the discrete-time samples $\mathbf{r}[n]$ of the wavefield observed by an array:

$$\mathbf{r}[\mathbf{n}] = \begin{bmatrix} r(\mathbf{x}_1, n\Delta t) \\ r(\mathbf{x}_2, n\Delta t) \\ \vdots \\ r(\mathbf{x}_M, n\Delta t) \end{bmatrix} \quad (34)$$

We treat the observed array signal as a vector process: quantities indicated in bold lower case are vectors and in bold upper case, matrices. Scalar quantities are indicated with an italic font.

The vectors $\{\mathbf{x}_i\}$ represent the locations of the sensors. M is the number of sensors in the array, and each of the i sensors observes the ground motion sequence $r(\mathbf{x}_i, n\Delta t)$. The sampling interval is denoted by Δt .

We take a quasi-frequency domain approach in which the data are partitioned into a large number (N) of narrow frequency bands through a bank of bandpass filters:

$$r_l[\mathbf{n}] = \sum_k h_l[k] r[\mathbf{n} - k] \quad (35)$$

The frequency-domain response of each of the filters is approximately one in a band centered at frequency $l\Delta f$ with bandwidth $\Delta f = 1/N\Delta t$ and is zero at all other frequencies. The impulse responses of the filters are related to a baseband lowpass filter $h_0[k]$ by complex exponential modulations:

$$h_l[k] = e^{i\Omega_0 l k} h_0[k]; l = 0, \dots, N-1; \Omega_0 = 2\pi\Delta f\Delta t \quad (36)$$

Defining as $H_0(\Omega)$ the frequency response (discrete Fourier transform) of the baseband filter:

$$H_0(\Omega) = \sum_k h_0[k] e^{i\Omega k} \quad (37)$$

then the choice of (3) for the family of filterbank impulse responses produces a collection of filters that are frequency translations of the baseband filter:

$$H_1(\Omega) = H_0(\Omega - \Omega_0) \quad (38)$$

It is possible to choose the baseband impulse response in such a way that the summed response of the filters is 1 across the frequency band of interest:

$$\sum_i H_i(\Omega) = 1 \quad (39)$$

which allows exact reconstruction of the original signal from its narrowband components. The filters $H_i(\Omega)$ are one sided in the frequency domain, hence complex. Consequently, the narrowband waveforms $r_i[n]$ are complex analytic signals.

In the time domain the beamforming operation consists of shifting operations on each waveform channel followed by summation of the results. In the frequency domain, the operation corresponding to a shift is multiplication of the Fourier transform of the i^{th} waveform by a frequency-dependent complex exponential phase factor

$$e^{-i\omega\Delta_i} \quad (40)$$

In the plane wave model of propagation, the delay $\Delta_i = \mathbf{v} \cdot \mathbf{x}_i$, where \mathbf{v} is the slowness vector defining (reciprocal) speed and the direction of propagation. The narrowband approximation to the frequency-domain phase shifting product consists of multiplication of each channels complex analytic representation by a phase factor:

$$e^{-i2\pi f\Delta_i} \quad (41)$$

which is constant in band Δf . The validity of this approximation depends on the bandwidth Δf of the filters being sufficiently small. Sufficiently small is defined by

$$\frac{\Delta f \Delta_i}{2} \ll 1; \forall i \quad (42)$$

so that phase changes over the width of each filter band are insignificant. With this condition, it can be shown that the beam is well-approximated by:

$$\mathbf{b}[n] = \sum_l \mathbf{w}_l^H \mathbf{r}_l[n] \quad (43)$$

where the complex weight vector \mathbf{w}_l assembles the phase delay factors for all channels:

$$\mathbf{w}_l = M^{-1/2} \begin{bmatrix} e^{-i2\pi l \Delta f \cdot x_1} \\ e^{-i2\pi l \Delta f \cdot x_2} \\ \vdots \\ e^{-i2\pi l \Delta f \cdot x_M} \end{bmatrix} \quad (44)$$

The scale factor $M^{-1/2}$ is chosen to assure that $\mathbf{w}_l^H \mathbf{w}_l = 1$. The vector \mathbf{w}_l is known as the steering vector for band l , and the inner product $\mathbf{w}_l^H \mathbf{r}_l[n]$ in (43) implements the narrowband equivalent of the shift and sum operation in that band.

Matched field processing generalizes beamforming by replacing the plane-wave steering vectors of beamforming with steering vectors estimated from prior observations of events at a target location. The estimation procedure consists of performing the narrowband decomposition of the prior observations, then computing covariance matrices for the vector waveforms in each band:

$$\mathbf{R}_l = \sum_n \mathbf{r}_l[n] \mathbf{r}_l^H[n] \quad (45)$$

If more than one event is available as a prior observation, these covariance matrices can be averaged over events. The steering vector for band l is obtained as the top eigenvector in the eigendecomposition:

$$\mathbf{R}_l = \mathbf{E}_l \mathbf{\Lambda}_l \mathbf{E}_l^H \quad (46)$$

Assuming $\mathbf{\varepsilon}_l$ is the eigenvector corresponding to the largest eigenvalue $(\lambda_l)_{\max}$, then we choose

$$\mathbf{w}_l = \mathbf{\varepsilon}_l \quad (47)$$

In the event that the design data consist of a noiseless observation of a pure plane wave, this procedure will reproduce the beamforming steering vector of equation (44). However, if the

observed signal departs significantly from a plane wave arriving along the great circle path from source to array, the principal eigenvector or eigenvectors will capture the scattering and refraction details not modeled by a plane wave.

We are examining two options for constructing a statistic $\gamma[n]$ for purposes of declaring detections. The first is an incoherent stack over the narrow bands:

$$\gamma_i[n] = \sum_i |w_i^H r_i[n]|^2 \quad (48)$$

This statistic is the classical wideband detection statistic of matched field processing from the underwater sound community. It estimates the power separately in each band and stacks these estimates incoherently over frequency. Its advantage is that it does not require the steering vectors to correctly preserve phase in the beam across frequencies. Its disadvantage is that it has very poor temporal resolution, roughly the resolution of the impulse response of a single narrowband filter, approximately $1/\Delta f$ seconds.

The second option is a coherent detection statistic, consisting of the squared magnitude of the wideband beam (equation 43):

$$\gamma_c[n] = |b[n]|^2 = \left| \sum_i w_i^H r_i[n] \right|^2 \quad (49)$$

The advantage of this detection statistic is that it potentially has the temporal resolution of a conventional beam. Its disadvantage is that it does require the steering vectors to preserve phase in the beam across the array. In principle, this requirement may be met by referencing the beam to one particular sensor by normalizing the steering vectors to have purely real components for that sensor. In practice, normalization may be hard to achieve if the signal is low or absent in a particular band for that sensor.

We digress at this point to give an example of steering vector calibration. The setting for the example is shown in Figure 25. We examine the use of one Myanmar event (2007 5/16 8:56:16.0, M_w 6.3; 20.47°N 100.69°E) as a calibration for a second event occurring nearly 4 years later (2011 3/24 13:55:12, M_w 6.9; 20.69°N 99.82°E). Both events are large and were recorded by the ILAR array with high SNR at many frequencies (see waveforms in Figure 26). We followed the design procedure just outlined for estimating steering vectors from waveforms of the 2007 event, using $N = 512$ bands. Since the data are real, only 257 of these bands are independent,

and it suffices to estimate steering vector in just the bands with indices $l = 0, \dots, 256$. The bandwidth for the narrowband filters in this analysis is $20/512 = 0.003904$ Hz (the data are sampled at 20 sps). The covariance matrices were computed over windows containing the first 25 seconds of the P arrival, for both the design event and the later target event.




Figure 25. Teleseismic paths from earthquakes in Myanmar to three North American arrays. The path length to ILAR (the nearest array) is about 8950 kilometers. Tohoku aftershock cloud shown in yellow.




Figure 26. Waveforms of Myanmar calibration event (left) and target event (right), recorded at ILAR. Waveforms from nine stations (IL01-IL09) are shown. Note that the second event is closely preceded by a smaller event, perhaps a Tohoku aftershock. The signals have been filtered into the 1-3 Hz band.




Figure 27. Comparison of empirical matched field processing (black) and plane wave (red) power for the 2011 3/24 M_w 6.9 event.

A comparison of beamforming performance for plane wave and empirical matched field steering vectors is shown in Figure 27. The measurements shown in this figure is the normalized power:

$$\frac{\mathbf{w}_1^H \mathbf{R}_1 \mathbf{w}_1}{\text{tr}\{\mathbf{R}_1\}} \quad (50)$$

The plane-wave steering vectors (equation 44) were determined by the great-circle backazimuth (299.2 degrees) from the array to the source and by the P velocity predicted by IASP-91 for the epicentral distance 8950 km. We see that below 2 Hz, the plane wave and the empirical steering vectors are performing nearly equally well – sometimes one is better, sometimes the other; values are comparable. Except in the vicinity of 1 Hz, where the empirical steering vectors appear to have a distinct advantage. Above 2 Hz, the empirical steering vectors perform substantially better than the theoretical plane wave vectors. This observation is consistent with expectations that scattering and refraction – driven by heterogeneities in the propagation medium – should increase with increasing frequency.

Beamforming and matched field processing are optimum estimators of the signal (and optimum detectors) in the case that the background noise is white, uncorrelated among sensors and of equal power on all sensors. In the event that the noise is correlated among sensors, the

optimum choice for steering vectors is the adaptive beamformer weighting [Capon et al., 1967] also known as the minimum variance, distortionless receiver (MVDR):

$$\mathbf{w}_l = \frac{\mathbf{R}_{nl}^{-1} \mathbf{e}_l}{\mathbf{e}_l^H \mathbf{R}_{nl}^{-1} \mathbf{e}_l} \quad (51)$$

To obtain these weights, an estimate of the noise covariance \mathbf{R}_{nl} is required, which means that a separate sample of noise data must be available for estimating covariance. In the example we show later, the “noise” will, in reality be non-stationary aftershock observations. The theory of adaptive beamforming was developed for stationary processes, but our example shows that it works to a degree for strongly non-stationary processes such as interfering aftershocks. In any case, the MVDR weighting has the effect of altering the wavenumber response of the array to suppress wavenumbers corresponding to strong noise components, while passing a desired signal with spatial structure defined by the steering vector \mathbf{e}_l .

2.3.2 Application: Detection of Signals of Monitoring Interest During an Ongoing Aftershock Sequence

In this section, we give an example of plane wave and matched field processing for the incoherent case (equation 48), both for the “conventional” (47) and adaptive (51) processors. The objective is to demonstrate improvements in beamforming performance in a spotlight situation during a major aftershock sequence. The example is contrived by burying the 2006 North Korean nuclear test among aftershocks of the 2008 Sichuan (Wenchuan) earthquake at a 1:1 ratio using data from the 19-channel ASAR (Alice Springs) array. Figure 28 shows the geographic configuration of the example. Figure 29 shows two channels of the superimposed ASAR data (sampling frequency: 20 sps). At its natural amplitude, the signal from the test is much smaller than many of the aftershock signals.




Figure 28. Geographic configuration of the test with superimposed Wenchuan aftershocks and the 2006 DPRK nuclear test. The North Korean test site is shown with the black outlined star and the observing array (ASAR) is shown with the red outlined star. The Wenchuan sequence aftershocks for the first day (May 12, 2008) are shown in white.

The results of processing the array data with four different matched field processing algorithms is shown in Figure 30. The figure shows detection statistics for 19,000 seconds of data, starting 6,000 seconds from a reference point (the main Sichuan event); signals from the North Korean test were buried approximately 18,000 seconds into the record from the main shock.




Figure 29. 2006 DPRK test superimposed at natural amplitude among aftershocks of the 2008 Wenchuan sequence. Waveforms from two channels (AS01, AS10) of the 19-channel ASAR array are depicted. The signal from the test appears at 6000 seconds in this plot. The data have been filtered into the 1-3 Hz band.

The top trace in Figure 30 displays the power output (equation 48) of an empirical matched field processor operating in the 1-3 Hz band (all four examples are in this band and all traces are plotted to the same scale). Waveforms from the 2009 North Korean tests observed at ASAR were used to obtain the signal model. The second trace shows the power output of a matched field processor which used the best plane-wave model for the wavefield incident from the North Korean test site. Comparing these two traces it is apparent that the empirical processor captures about twice the power from the desired signal as that captured by the plane wave processor. The empirical processor also suppresses the aftershocks slightly better than the plane wave processor.

The bottom two traces show the corresponding power traces for the empirical (third trace) and plane-wave (fourth trace) MVDR matched field processors. MVDR processors require an estimate of the background noise covariance, which in this case was supplied by aftershocks in the first 6,000 seconds of data (prior to the displayed traces). Both processors have done an impressive job suppressing aftershock power, but the empirical processor again has captured

about twice the energy from the North Korean test compared to the plane-wave processor. Although it is less apparent, the empirical processor has done substantially more to suppress aftershock power than the plane-wave processor with reference to the non-MVDR case.




Figure 30. Detection statistics for four matched field processors operating on data from the ASAR array show that empirical matched field processing with MVDR sidelobe suppression improves detectability of a weak signal among strong aftershocks from a different location.

We anticipate that more dramatic improvements (comparing empirical MVDR to plane-wave MVDR matched field processing) will be obtained with regional observations in higher frequency bands than shown here. The effects of heterogeneities on the signal model are more severe at regional distances.

2.3.3 Application: Detection of Signals of Monitoring Interest in the Coda of a Large Earthquake

In this section, we give a second example of matched field processing for the incoherent (equation 48) and coherent (equation 49) cases, both for the “conventional” (47) and adaptive (51) processors. In this scenario, we consider detection of the 2016 DPRK explosion in the coda of the second Nepal earthquake of May 12, 2015 (Mw 7.3, 12:50:19, 27.837N 86.077E) at the Alaskan array ILAR using waveforms of the 2013 DPRK explosion as the template (Figure 31). Our objective in constructing this scenario is to demonstrate the superiority of adaptive beamforming in detecting an event of monitoring interest for which signals are masked by arrivals from a larger earthquake.




Figure 31. Great circle paths from the DPRK nuclear test site (red) and the Nepal earthquake sequence (white) to the ILAR array in Alaska.

Figure 32 shows, on a large scale, the construction of the test waveforms obtained by embedding the ILAR observations of the 6 January 2016 DPRK nuclear test in the coda of the second Nepal earthquake. Although nineteen ILAR channels were used to create the scenario data stream, only IL02 is shown in the figure. All data were processed in a teleseismic frequency band (1.5-3.5 Hz).




Figure 32. Scenario waveforms obtained by superposition of signals (S1) recorded at ILAR for the 12 May 2015 Nepal earthquake and the P wave observations (S2) also recorded at ILAR for the 6 January 2016 DPRK nuclear test. The signals are superimposed at a 1:1 amplitude ratio (true amplitude).

Detail of the superposition is shown in Figure 33. In this plot, the third trace shows that the embedded explosion P wave waveform, while visible, is not distinguishable by eye from scattered energy bursts in the coda (it presumably would be distinct in a plot of the distribution of waveform energy in time and azimuth).




Figure 33 Detailed plot of the superposition of the target waveform ($S2$) in coda of the Nepal earthquake ($S1$). The embedded explosion P wave is unremarkable, being similar in size to bursts of scattered energy from the earthquake.

The next figure (34) shows the power output from two frequency-incoherent matched field processors (equation 48) using the conventional (47) steering vector (black in the figure) and the adaptive weights (51, red in the figure). The conventional beam power is dominated by earthquake P and P coda; a small peak corresponding to the explosion P wave exists, but is indistinguishable from secondary peaks due to the earthquake coda. The MVDR weights in this example were constructed from a “noise” covariance matrix estimated from the entire data window of the superimposed signals (third trace of Figure 32). Because the signals from the Nepal earthquake are so much more energetic than the explosion P wave, the covariance estimate is dominated by the P and P coda of the earthquake, the portion of the stream that we wish to suppress.

The red trace of Figure 34 demonstrates that adaptive beamforming does indeed effectively suppress the earthquake wavetrain. The bottom half of the figure compares the two power traces expanded vertically and cropped to show greater detail. In this bottom part of the figure, the adaptive processor power trace has a clear peak due to the explosion P phase. This peak is

now the largest feature of the trace, indicating that adaptive processing is effective in allowing the smaller event to be observed.




Figure 34. Power traces (with square root scaling) for frequency-incoherent matched field processors. The black trace shows the power output by a "conventional" processor and the red trace shows power output by the adaptive processor. A peak corresponding to the embedded P phase from the DPRK event is clear in the adaptive beam power trace, but is a minor feature in the earthquake coda in the conventional beam power trace.

Not all of the energy from the earthquake arrivals is suppressed, leaving a pedestal on which the peak from the target signal is superimposed. But enough is removed to allow the presence of an event of interest to be inferred.

Figure 35 compares power output traces for the frequency-coherent (49) matched field counterparts of incoherent processors just discussed. The coherent processors are the direct matched field equivalents of shift and sum beamforming. The result is quite similar to the incoherent case: the conventional matched field processor power does not show a

distinguishable peak identifying the embedded explosion signal, but the adaptive processor does. The coherent processor results are remarkable for two characteristics: resolution of peaks in the time domain is improved, and the SNR of the target peak is increased.




Figure 35. Power traces (with square root scaling) for frequency-coherent matched field processors also demonstrate that adaptive processing is effective in suppressing unwanted interference, allowing the existence of a target signal to be inferred. Here, as in the last figure, the black trace represents the power output of a “conventional” matched field processor and the red trace represents the power output of an adaptive processor. A clear peak in the red trace in the bottom expanded view marks the arrival of the embedded explosion P wave.

The slightly increased performance of the coherent processor comes at the expense of greater algorithmic complexity, but a fast algorithm for synthesizing the wideband beam from its narrowband components is available, making this approach computationally attractive.

2.3.4 Application: Detection of Signals of Monitoring Interest Among Mining Explosions and Earthquakes

As our final example of adaptive beamforming, we examine the problem of detecting a nuclear explosion in a heterogeneous assemblage of regional and teleseismic events with conventional and adaptive matched field processors. One channel of the data stream we examine is shown in Figure 36. The stream consists of 12 channels from the Wyoming array PDAR, beginning at 2015:132:11:00:00.000 and continuing for 51,400 seconds. During this period, the array recorded mining explosions from the Powder River Basin and at least one teleseism (the second large signal about $\frac{3}{4}$ of the way through the record). The waveforms from the 12 February 2013 DPRK explosion were embedded in this record about 12,000 seconds from the start (making it the second obvious event), at a 1:1 amplitude ratio. The objective of this test was to examine how far adaptive processing suppresses regional signals when searching for a target teleseism.




Figure 36. Superposition of continuous stream data from station PD01 and the 2013 DPRK explosion. The target DPRK event is the second obvious event about 12,000 seconds into the record.

In this case, the 6 January 2016 DPRK test waveforms recorded at PDAR were used to provide the multichannel template for matched field processing. The results of processing the data with incoherent matched field processors in a teleseismic band (1-3 Hz) is shown in Figure 37. The top trace is output from the conventional matched field processor and the bottom is output from the adaptive processor. In this test, the adaptation took place over the interval from 5,000 seconds to 45,000 seconds, i.e. the noise covariance estimate was made over almost the entire data interval. The resulting adaptive processor is more effective at suppressing the regional and

teleseismic events in the record, but the results are not as dramatic as the suppression of the P and P coda arrivals from the Nepal earthquake at ILAR. This outcome probably arises from two factors: the collection of events is far more heterogeneous than the interference in the last example, and the array has fewer elements (12 vs. 19).




Figure 37. Frequency-incoherent power traces (with square-root scaling) for conventional (top) and adaptive (bottom) matched field processors. The adaptive processor more effectively suppresses competing events.

The array has fewer degrees of freedom to throw at the interference suppression problem and the events exhibit more wavenumber spectrum diversity. Interference energy in the ILAR example of the last section was dominated by the main shock of the Nepal sequence; most subsequent events were aftershocks. Consequently, the events present in that example occupied a much smaller patch of wavenumber space, allowing the array response to be adjusted by the adaptive processor to suppress that small region more effectively. This observation suggests that a processor that adapts continuously over much shorter time horizons would be required to obtain better suppression of diverse interference. In addition,

arrays with larger numbers of elements will have the degrees of freedom to provide better interference suppression.

In the last figure (38), the amplitude envelope of a single array channel is compared to the square root of the beam power trace for an adaptive coherent processor. This comparison gives an idea of the overall gain of the best array processing. By comparison with Figure 37, we see that performance of the coherent processor is only marginally better than that of the incoherent processor in this case.




Figure 38. Amplitude envelope of a single raw data channel (top) compared to the power output (with square-root scaling) of the frequency-coherent matched field processor (bottom). This figure indicates overall processing gain. The two traces have been scaled so that the peak due to the DPRK explosion (second from left) has approximately the same amplitude in both plots.

2.3.5 Summary

We sought to improve the sensitivity of signal detectors using empirical matched field processing (EMFP) and adaptive beamforming. Matched Field Processing can be considered to be a form of frequency domain beamforming and EMFP is a calibrated form where the applied phase shifts are calculated by the observations of previous signals from the direction of interest. EMFP provides the optimal estimator of the signal under the assumption of white background noise. In the case of correlated noise, the optimal choice for steering vectors is the adaptive beamformer weighting – or minimum variance distortionless receiver (MVDR).

To demonstrate, signals from confirmed nuclear tests in the DPRK were submerged into aftershock sequences recorded on medium aperture seismic arrays (with aperture of the order 10 km) with the noise covariance structure for MVDR estimated from a time average of the ongoing aftershock sequence. Relative to the standard matched field estimator, adaptive processors reduce, often dramatically, the contribution of interfering events to detection statistics traces targeted on the DPRK test.

We draw two conclusions from the examples provided in this section: (1) the use of calibrated steering vectors with empirical matched field processing appears to allow adaptive beamforming techniques (heretofore largely abandoned in seismology) to be used successfully, and (2) adaptive beamforming may be used to advantage in several stages of iteration in an advanced pipeline. It may support more sensitive initial event detection in the pipeline – the example of section 2.3.2 was of this type. In this case the adaptive beamforming weights were constructed from observations of interference prior to the occurrence of the signal of interest. In addition, adaptive beamforming may support context-driven reprocessing of a stream at a station which failed to make a detection when observations at other stations in a network indicate that it should have had a trigger. The example of section 2.3.3 was of this type; in this case, the adaptive weights were constructed from interference correlation measurements made in a window that included the target signal.

2.4 Evaluation of Event Hypotheses

A context-driven processing framework will attempt to evaluate how consistent candidate event hypotheses are with the observations. Subsequent reprocessing of the data stream(s) will depend upon the event hypotheses which are considered by the previous iteration to best represent the true event history. For this we need an empirical (or semi-empirical) means of predicting the likely observations from a given hypothetical event history. We need realistic estimates of the likelihood of a set of observations from a given event hypothesis.

An event hypothesis may be based on a small number of arrivals which, although technically consistent with an event hypothesis at a given location and a given time, would not be expected to be the only observations seen. If an event hypothesis is missing detections from key stations which would be expected to detect signals, and yet includes phases at stations for which the likelihood of detection is very low, the hypothesis likely to be incorrect. To address this issue, NORSAR has over some time assessed detection capabilities for IMS stations for application to the global phase association process (Kværna and Ringdal, 2013). The basis for this work is the events and phase readings reported in the IDC reviewed Event Bulletin (REB), as well as the corresponding lists of non-detecting stations. One type of result from this study is a map of regionalized capability estimates for each IMS station (see Figure 39). To provide station detection capability estimates for regions with no, or very few, calibration events, a standard amplitude-distance curve (e.g. Murphy and Barker, 2003) is fitted to regionalized estimates.

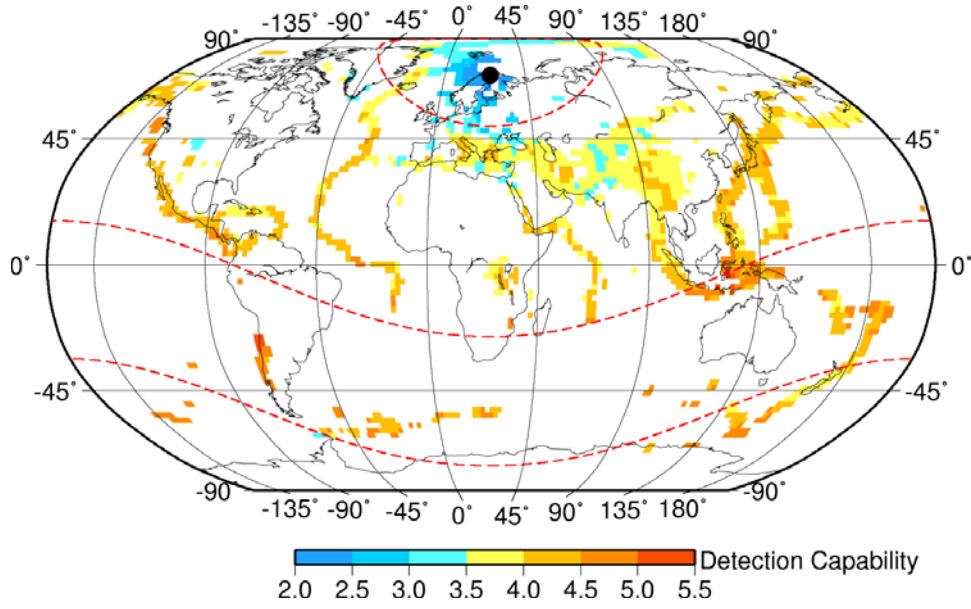


Figure 39. *P*-phase detection capability estimates for the ARCES array in 2°x2° bins. The red stippled curves denote distances of 20, 95 and 144 degrees from ARCES. The blue bins denote “bright spots” where the ARCES detection capability is particularly good ($m_b \leq 3.5$).

2.4.1 Estimation of Detection Thresholds for Specific Source Regions

Taking the North Korea test site as an example, we can, based on several years of REB bulletin data, estimate the likely detection capability at a given set of stations using the approach of Kværna and Ringdal (2013). Figure 40 displays the 44 best-case detection capability estimates for this test site for IMS stations (with arrays and 3-component stations differentiated) and Figure 41 displays the detection capability estimates for this set of stations as a function of epicentral distance. We find that these station detection capability estimates are in good agreement with the distribution of signal-to-noise ratios reported in the REB for the North Korean underground tests.




Figure 40. Predicted detection thresholds in units of magnitudes for the 44 IMS stations (both array and 3-component, primary and auxiliary) with best predicted detection capability for the North Korea nuclear test site. (Assumed location 41.28°N, 129.07°E). Array stations (both primary and auxiliary) are displayed in red and 3-component stations (both primary and auxiliary) are displayed in blue.




Figure 41. Predicted detection thresholds in units of magnitudes for the North Korea nuclear test site (assumed location 41.28oN, 129.07oE) as a function of epicentral distance for the 44 IMS stations displayed in Figure 3. The array stations are displayed using red circles and the 3-component stations are displayed using blue triangles.

Any event hypothesis resulting from a global association of phases from the parametric datastream should show a certain consistency with these two displays. This is to say that detection of an event at a station with low detection probability and non-detection at stations with high event probability should immediately flag an event hypothesis as questionable. In the design of optimal beams/detection statistic traces to evaluate the validity of event hypotheses, such detection capability plots should form the basis of all detectors spawned for hypothesis confirmation.

The locations of the stations displayed in the previous two figures are shown in Figure 42.

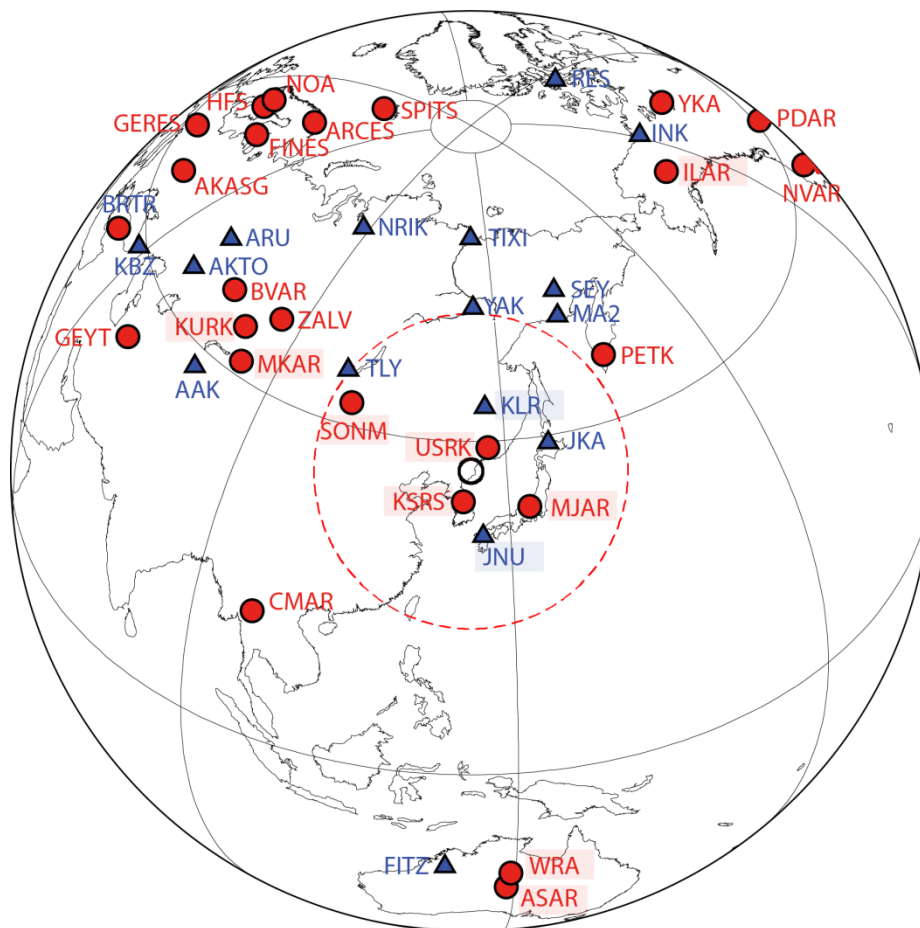


Figure 42. Locations of the stations displayed in the previous two figures. The most sensitive stations for this source region have the station codes shaded.

The primary goal of this study is to assess the consistency of the automatic phase associations of the IDC SEL's, and to improve the performance of the phase association algorithm. However, the information about the station detection capabilities, as shown in Figure 40, can equally well be used as a tool for recall data processing of stations having a high likelihood of observing signals from a hypothesized event. Similarly, recall data processing can be done to reject incorrectly associated phases at stations having a low probability of detecting the event.

It is the intention to exploit this information in several ways:

- 1) to predict the stations that should be severely affected by a developing aftershock sequence,
- 2) to identify the stations that should observe certain phases for an hypothesized event, but failed to do so, and

3) to identify situations where stations reported defining phase detections (those used to build an event), but were unlikely to observe phases for the hypothesized event.

These are all situations where supervisory function in a future pipeline might initiate tailored processing or reprocessing of waveform data. In the first case, beamforming algorithms might be replaced for the duration of the aftershock sequence with adaptive methods suitable for suppressing aftershock waveforms. In the second case, specialized beams might be deployed to search for the missing phases. And in the third case, specialized beams might be used to confirm or refute the existence of defining phases at the reporting stations. This last process would serve as a check on improperly associated phase detections. In all three cases, adaptive beamforming algorithms might be used to improve on results obtained by conventional (recipe) beamforming algorithms in a first pass over the data.

A context-driven processing framework will attempt to evaluate how consistent candidate event hypotheses are with the observations. Subsequent reprocessing of the data stream(s) will depend upon the event hypotheses which are considered by the previous iteration to best represent the true event history. For this we need an empirical (or semi-empirical) means of predicting the likely observations from a given hypothetical event history.

2.4.2 Evaluating Automatically Generated Event Hypotheses

A significant component of the design of iterative processing pipelines is to evaluate approaches for accepting or rejecting the association of phases that contribute to a given event hypothesis. Comparing the (fully automatic) SL3 event bulletin and the (reviewed) REB from the International Data Center (IDC) of the Comprehensive Nuclear-Test-Ban Treaty Organization (CTBTO) indicates a number of circumstances in which false event hypotheses have been formed from genuine phase detections and true events which have had to be built manually.

Figure 43 shows an example of such a case. Two events in central Asia have generated waveforms at stations globally that overlap. The result is that the first event, a deep earthquake in the Hindu Kush region of Afghanistan, is well classified automatically with a qualitatively correct solution in the automatic SL3 bulletin. The second event, a shallow earthquake close to the epicenter of the M=7.6 October 8, 2005, Kashmir event, generates signals which are associated with a spurious event hypothesis.

```

EVENT 3441778 HINDU KUSH REGION, AFGHANISTAN
  Date      Time      Err    RMS Latitude Longitude  Smaj  Smin  Az Depth  Err  Ndef  Nsta  Gap  ndist  Mdlist  Qual  Author  OrigID
2005/11/05 23:24:27.35 6.75 0.92 36.8415 70.7943 62.7 21.9 176 257.8 54.7 13 13 231 13.11 80.94 m i uk IDC_SEL3: 3441904

Magnitude Err Nsta Author OrigID
mb 3.5 0.2 10 IDC_SEL3: 3441904

Sta Dist EvAz Phase Time TRes Azim AzRes Slow SRes Def SNR Amp Per Qual Magnitude ArrID
MKAR 13.11 37.1 P 23:27:23.175 -1.2 223.4 -1.4 17.1 4.8 TAS 53.9 0.9 0.33 a__ 26110551
BVAR 16.19 359.1 P 23:28:00.437 1.0 158.8 -20.1 12.4 1.5 TAS 49.8 2.6 0.33 a__ 26112040
AKTO 16.40 329.9 P 23:28:02.375 0.5 137.6 -3.5 10.9 -0.0 TAS 33.7 1.6 0.33 a__ 26111681
SONM 28.28 55.9 P 23:29:57.560 0.5 270.3 10.0 9.1 0.2 TAS 10.3 1.0 0.59 a__ mb 3.3 26110501
BRTR 29.17 287.1 P 23:30:04.550 -0.5 100.8 17.1 8.2 -0.6 TAS 8.5 1.1 0.63 a__ mb 3.3 26110565
AKASG 32.60 308.5 P 23:30:33.900 -0.8 80.6 -18.4 7.8 -0.8 TAS 6.7 0.4 0.35 a__ mb 3.1 26110563
MLR 34.47 298.8 P 23:30:53.384 2.4 169.4 81.0 3.3 -5.2 TAS 4.9 2.4 0.88 a__ mb 3.4 26111758
FINES 37.10 325.9 P 23:31:12.700 -0.2 114.6 -3.8 10.6 0.0 TAS 12.0 2.0 0.67 a__ mb 3.5 26110539
ARCES 40.76 337.5 P 23:31:43.250 0.3 111.7 -1.6 7.5 -0.2 TAS 29.2 2.6 0.44 a__ mb 3.8 26110680
HFS 42.70 321.7 P 23:31:57.994 -0.7 97.3 0.7 10.2 2.2 TAS 15.5 1.8 0.37 a__ mb 3.7 26111908
NOA 44.01 322.8 P 23:32:08.350 -0.8 96.5 0.2 7.7 -0.0 TAS 7.6 1.6 0.61 a__ mb 3.4 26110692
INK 73.61 9.2 P 23:35:32.375 0.4 328.7 -11.1 4.9 -0.9 TAS 11.8 1.8 0.61 a__ mb 3.7 26112063
YKA 80.94 2.5 P 23:36:11.945 -0.6 348.9 -3.4 5.4 0.0 TAS 6.0 0.6 0.34 a__ mb 3.5 26110683

EVENT 3441926 LAKE ISSYK-KUL REGION
  Date      Time      Err    RMS Latitude Longitude  Smaj  Smin  Az Depth  Err  Ndef  Nsta  Gap  ndist  Mdlist  Qual  Author  OrigID
2005/11/05 23:25:32.57 2.67 0.11 42.6983 79.7545 61.0 33.1 68 0.0f 3 3 203 4.48 80.05 m i uk IDC_SEL3: 3441929

Magnitude Err Nsta Author OrigID
ML 2.5 0.4 2 IDC_SEL3: 3441929
mb 4.1 1 IDC_SEL3: 3441929

Sta Dist EvAz Phase Time TRes Azim AzRes Slow SRes Def SNR Amp Per Qual Magnitude ArrID
MKAR 4.48 22.9 Pn 23:26:43.725 -0.1 189.5 -15.2 14.8 1.0 TAS 4.6 0.02 0.33 a__ ML 1.4 26110550
BVAR 12.08 332.0 Pn 23:28:24.062 -0.1 167.5 22.4 13.7 -0.0 TAS 3.9 0.4 0.33 a__ ML 3.6 26112161
WRA 80.05 128.9 P 23:37:44.259 0.1 320.4 -10.8 5.2 0.2 TAS 6.6 1.3 0.79 a__ mb 4.1 26110688
WRA 80.05 128.9 tx 23:37:57.384 320.4 -10.9 5.0 --- 7.9 0.5 0.68 a__ 26110689
WRA 80.05 128.9 tx 23:38:02.559 321.3 -10.0 5.0 --- 7.9 0.6 0.63 a__ 26110690

```

Figure 43. Extract from SL3 automatic event bulletin from the International Data Center (IDC) from November 5, 2005.

The analyst-reviewed location estimates (Figure 44) indicate that the Kashmir event has been reconstructed manually entirely from the Warramunga P-detection (no other detections in the SL3 event hypothesis are now associated with the REB event). The overlapping waveforms are displayed for a number of stations in Figure 45. It is demonstrated in Figure 46 how a fully automatic event location is possible by taking the WRA P-detection as a trigger and projecting this back over a source region and creating multiple event hypotheses and analyzing which event hypotheses can associate the greatest number of phases and the smallest time residuals. We note in this case that, even on optimal traces targeted towards the source region of the second event, accurate automatic determinations of the arrival times are exceptionally difficult. Automatic rejection of the corresponding SL3 event hypothesis is far simpler.

EVENT 3441778 HINDU KUSH REGION, AFGHANISTAN																
Date	Time	Err	RMS	Latitude	Longitude	Smaj	Smin	Az	Depth	Err	Ndef	Nsta	Gap	mdist	Mdist	Qual
2005/11/05	23:24:23.92	4.21	0.89	36.5181	70.8303	23.4	13.6	2	246.1	39.2	18	17	128	13.36	149.70	m i uk IDC_REB:I
OrigID 3443754																
Magnitude	Err	Nsta	Author	OrigID												
mb	3.5	0.1	12	IDC_REB:I	3443754											
mb1	3.6	0.1	16	IDC_REB:I	3443754											
mb1mx	3.5	0.1	23	IDC_REB:I	3443754											
mbtmp	4.2	0.1	16	IDC_REB:I	3443754											
Sta	Dist	EvAz	Phase	Time	TRes	Azin	AzRes	Slow	SRes	Def	SNR	Amp	Per	Qual	Magnitude	ArrID
MKAR	13.36	36.2	P	23:27:22.650	-1.6	223.4	-0.5	17.1	4.8	T__	53.9	0.9	0.33	a__	mb1	3.4 26110551
BVAR	16.51	359.1	P	23:27:59.762	-0.5	158.8	-20.0	12.4	1.4	T__	49.8	2.6	0.33	a__	mb1	3.8 26112040
BVAR	16.51	359.1	S	23:30:59.262	-1.5	172.4	-6.3	21.8	1.7	T__	5.0					26147621
AKTO	16.70	330.4	P	23:28:02.375	-0.0	137.6	-4.0	10.9	-0.0	T__	33.7	1.6	0.33	a__	mb1	3.9 26111681
ZAL	19.94	24.7	P	23:28:36.850	-0.4	248.8	-168.	6.0	-4.8	T__	19.7	2.7	0.33	a__	mb1	3.9 26110609
ZAL	19.94	24.7	S	23:32:04.175	-4.3	159.1	-78.9	21.1	2.4	T__	7.5	0.3	0.33	a__		26110694
BONH	28.44	55.4	P	23:29:57.560	1.5	270.3	10.6	9.1	0.3	T__	10.3	1.0	0.59	a__	mb	3.3 26110501
BRTR	29.30	287.7	P	23:30:04.550	0.8	100.8	16.5	8.2	-0.6	T__	8.5	1.1	0.63	a__	mb	3.3 26110565
AKASG	32.83	308.8	P	23:30:33.900	-0.3	80.6	-18.8	7.8	-0.8	T__	6.7	0.4	0.35	a__	mb	2.9 26110563
MLR	34.65	299.2	P	23:30:51.834	1.7	169.4	80.6	3.3	-5.2	T__	4.9	2.4	0.88	a__	mb	3.6 26111758
FINES	37.38	326.1	P	23:31:12.700	-0.1	114.6	-4.0	10.6	0.0	T__	12.0	2.0	0.67	a__	mb	3.5 26110539
ARCES	41.07	337.6	P	23:31:43.250	0.1	111.7	-1.8	7.5	-0.1	T__	29.2	2.6	0.44	a__	mb	3.8 26110680
HFS	42.97	321.9	P	23:31:57.994	-0.5	97.3	0.4	10.2	2.2	T__	15.5	1.8	0.37	a__	mb	3.7 26111908
NOA	44.28	323.0	P	23:32:08.250	-0.7	96.5	-0.0	7.7	-0.0	T__	7.6	1.6	0.61	a__	mb	3.5 26110692
INK	73.92	9.2	P	23:35:32.000	0.4	328.7	-11.0	4.9	-0.9	T__	11.8	1.8	0.61	a__	mb	3.7 26112063
DBIC	74.56	266.9	P	23:35:35.250	-1.0	213.3	178.2	18.0	12.3	T__	4.2	2.7	0.74	a__	mb	3.8 26110655
YKA	81.26	2.6	P	23:36:11.945	-0.1	348.9	-3.3	5.4	0.0	T__	6.0	0.6	0.34	a__	mb	3.4 26110683
WRA	82.12	121.8	P	23:36:16.409	-0.9	323.2	1.0	4.3	-0.4	T__	3.4	0.5	0.63	a__	mb	3.2 26147622
PLCA	149.70	250.1	PKPbc	23:43:44.400	0.4	152.9	67.5	8.2	5.7	T__	4.5	3.1	0.91	a__		26110706
EVENT 3441926 PAKISTAN																
Date	Time	Err	RMS	Latitude	Longitude	Smaj	Smin	Az	Depth	Err	Ndef	Nsta	Gap	mdist	Mdist	Qual
2005/11/05	23:25:36.11	1.52	0.56	34.2471	73.5121	36.3	24.2	43	0.0f		9	9	200	14.19	79.05	m i uk IDC_REB:I
OrigID 3443757																
Magnitude	Err	Nsta	Author	OrigID												
ML	4.2	0.0	2	IDC_REB:I	3443757											
mb	4.0	0.1	6	IDC_REB:I	3443757											
mb1	4.1	0.1	9	IDC_REB:I	3443757											
mb1mx	3.9	0.1	21	IDC_REB:I	3443757											
mbtmp	4.0	0.1	9	IDC_REB:I	3443757											
Sta	Dist	EvAz	Phase	Time	TRes	Azin	AzRes	Slow	SRes	Def	SNR	Amp	Per	Qual	Magnitude	ArrID
MKAR	14.19	25.3	Pn	23:28:58.999	-0.2	229.4	18.3	13.3	-0.3	T__	2.3				mb1	3.5 26147769
BVAR	18.90	354.2	P	23:29:58.265	0.2	164.1	-7.9	10.5	-0.7	T__	5.0	0.7	0.33	a__	ML	4.2 26147768
AKTO	19.75	329.6	P	23:30:08.525	1.0	145.2	6.1	11.1	0.2	T__	5.8	0.7	0.33	a__	ML	4.2 26111771
ZAL	21.24	18.6	P	23:30:22.550	-0.9	255.7	-167.	6.3	-4.5	T__	6.2	5.5	0.37	a__	mb	4.3 26110613
SONH	28.04	51.1	P	23:31:29.580	0.5	258.2	5.1	9.4	0.5	T__	4.7	0.8	0.79	a__	mb	3.7 26110622
ARCES	44.00	337.9	P	23:33:45.067	0.3	110.6	-1.2	7.7	0.2	T__	3.7	4.5	1.02	a__	mb	4.3 26147766
HFS	46.10	323.1	P	23:34:01.582	0.0	99.3	2.9	9.9	2.0	T__	9.5	1.0	0.34	a__	mb	4.1 26111909
NOA	47.41	324.2	P	23:34:11.172	-0.7	97.4	1.4	7.7	0.1	T__	2.9	0.8	0.55	a__	mb	3.8 26147767
WRA	79.05	123.2	P	23:37:42.409	-0.0	320.4	-1.0	5.2	0.2	T__	6.6	1.3	0.79	a__	mb	4.1 26110688
															mb1	4.5
															mbtmp	4.1

Figure 44. Extract from Reviewed Event Bulletin (REB) from the International Data Center (IDC). The first of these events was well predicted automatically (see the SL3 page displayed in Figure 43), and the second of these events was built manually from a single phase extracted from a false event hypothesis. The relocated events are separated by approximately 350 km, although the first event is in addition over 250 km deep.




Figure 45. Waveforms for selected stations (optimized for given phases from the Kashmir region) with arrivals from the events listed in Figure 44 labelled.




Figure 46. Source-scan for triggering WRA P-phase at time 2005-309:23.37.43. A geographical grid was searched with origin times determined by the theoretical traveltimes to WRA and the number of automatic phase detections from other stations consistent with each of these event hypotheses was counted. Of all of the event hypotheses with the maximum number of phases (8), the hypocenter with the minimum traveltime residual was 2005-309:23.25.40 at 34.46°N, 73.60°E and 20 km deep. The colors in this plot indicate the 1-norm traveltime residual for this set of phases over a far denser grid of trial hypocenters. The optimal epicenter is at 34.477°N 73.524°E.

This is a pathological example of completely unrelated overlapping events which are likely to provide a challenge to a context-based detection framework. These events will appear very closely in slowness space for almost all global arrays and may expose limiting factors for a number of the techniques proposed here.

2.5 Iterative Schemes for Automatic Event Bulletins: Coping with Aftershock Sequences

An obvious application of iterative pipeline processing is the event of large earthquakes followed by extensive aftershock sequences. The rapidly unfolding sequences of phase arrivals at the stations in a regional or global network can present significant challenges to automatic procedures for phase association and the creation of event hypotheses. On May 25, 2015, an earthquake of magnitude 7.8 struck Nepal and was followed by several hundreds of aftershocks which were recorded globally. Figure 47 displays the event hypotheses in the fully automatic SL3 bulletin of the International Data Center (IDC) in the first six days of this aftershock sequence, together with the final analyst-reviewed solutions for the same time interval. This figure illustrates at a glance the problem faced by automatic algorithms in constructing an accurate representation of the sequence of seismic events.

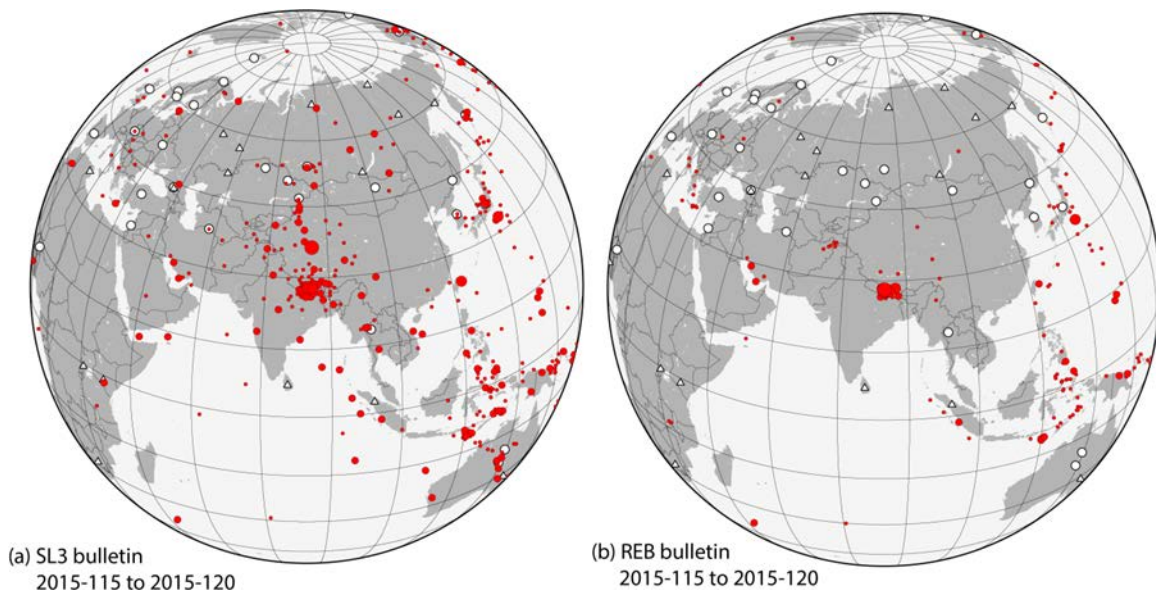


Figure 47. Event location estimates from the fully automatic SL3 bulletin from the International Data Center (IDC) in Vienna for April 25 to April 30, 2015, and (b) event locations from the analyst-reviewed Reviewed Event Bulletin (REB) for the same time period.

In the classical single-pass approach to processing pipelines, a sequence of detections with associated parametric data is forwarded from each station in the network to an association algorithm which attempts to form a preliminary automatic event bulletin. In the aftershock sequence scenario, we have issues such as the overlap of signals (where first arrivals at key stations may be buried in the coda of signals from earlier events) and a succession of events that is so rapid that qualitatively incorrect solutions with spuriously associated arrivals will often score more highly in automatic evaluations than the often incomplete sets of detections that

correspond to the true origin time and source location. This is what leads to a scatter of preliminary event location estimates over continental distance scales (Figure 47 a). It is the enormous human effort involved in producing the result displayed in Figure 47 b) from the starting point displayed in Figure 47 a) that we seek to assist. While there may be many strategies for improving the association algorithms (e.g. Arora et al., 2013), we wish to explore context-based iterative solutions where new passes over the raw waveform data are performed using techniques which are optimized for the understood state of seismicity.

The principles of existing processing pipelines which lead to the event distributions displayed in Figure 47 are illustrated in Figure 48. The datastreams (a) are present for many stations in a regional or global network. While only one stream is displayed here for each station, in practice many different detection statistic traces are formed for each station. For 3-component stations, detection processes can be run independently on different rotations of the orthogonal components to seek P- and S- phases from different directions, and processed in multiple frequency bands in order to optimize the signal-to-noise ratio (SNR) for a broad range of anticipated signals. For seismic arrays, the situation is even more complicated with an extensive beam deployment steering numerous beams with carefully selected combinations of parameters for optimizing the detection of signals from all anticipated seismic sources. For all stations, a detection reduction and parameter estimation process will result in parametric datastreams (Figure 48 b) whereby, for a given trigger time, we have a set of parameters describing the nature of the detection. These parameters include typically SNR, apparent velocity, backazimuth, frequency content, and amplitude. A Global Association algorithm will determine a set of event hypotheses which best explain the sets of parametric data, an automatic event bulletin (Figure 48 c). Four events have been formed from the parametric data streams in this illustrative example. Three events, the red, the blue and the green, come from essentially the same geographical location but are placed at some distance from each other due to inaccuracies in the parametric data and/or a deficiency in the phase association algorithm. This situation depicted here is what leads to the complicated real-life example shown in Figure 47 a). The analyst goes through the available sets of automatic bulletins, with access to the waveform data (e), and relocates events to form the reviewed bulletin (d). Event hypotheses that are clearly false are deleted and new events that were not identified by the automatic processing are added if the analyst sees sufficient evidence. Panels c) and d) of Figure 48 correspond to panels a) and b) of Figure 47.

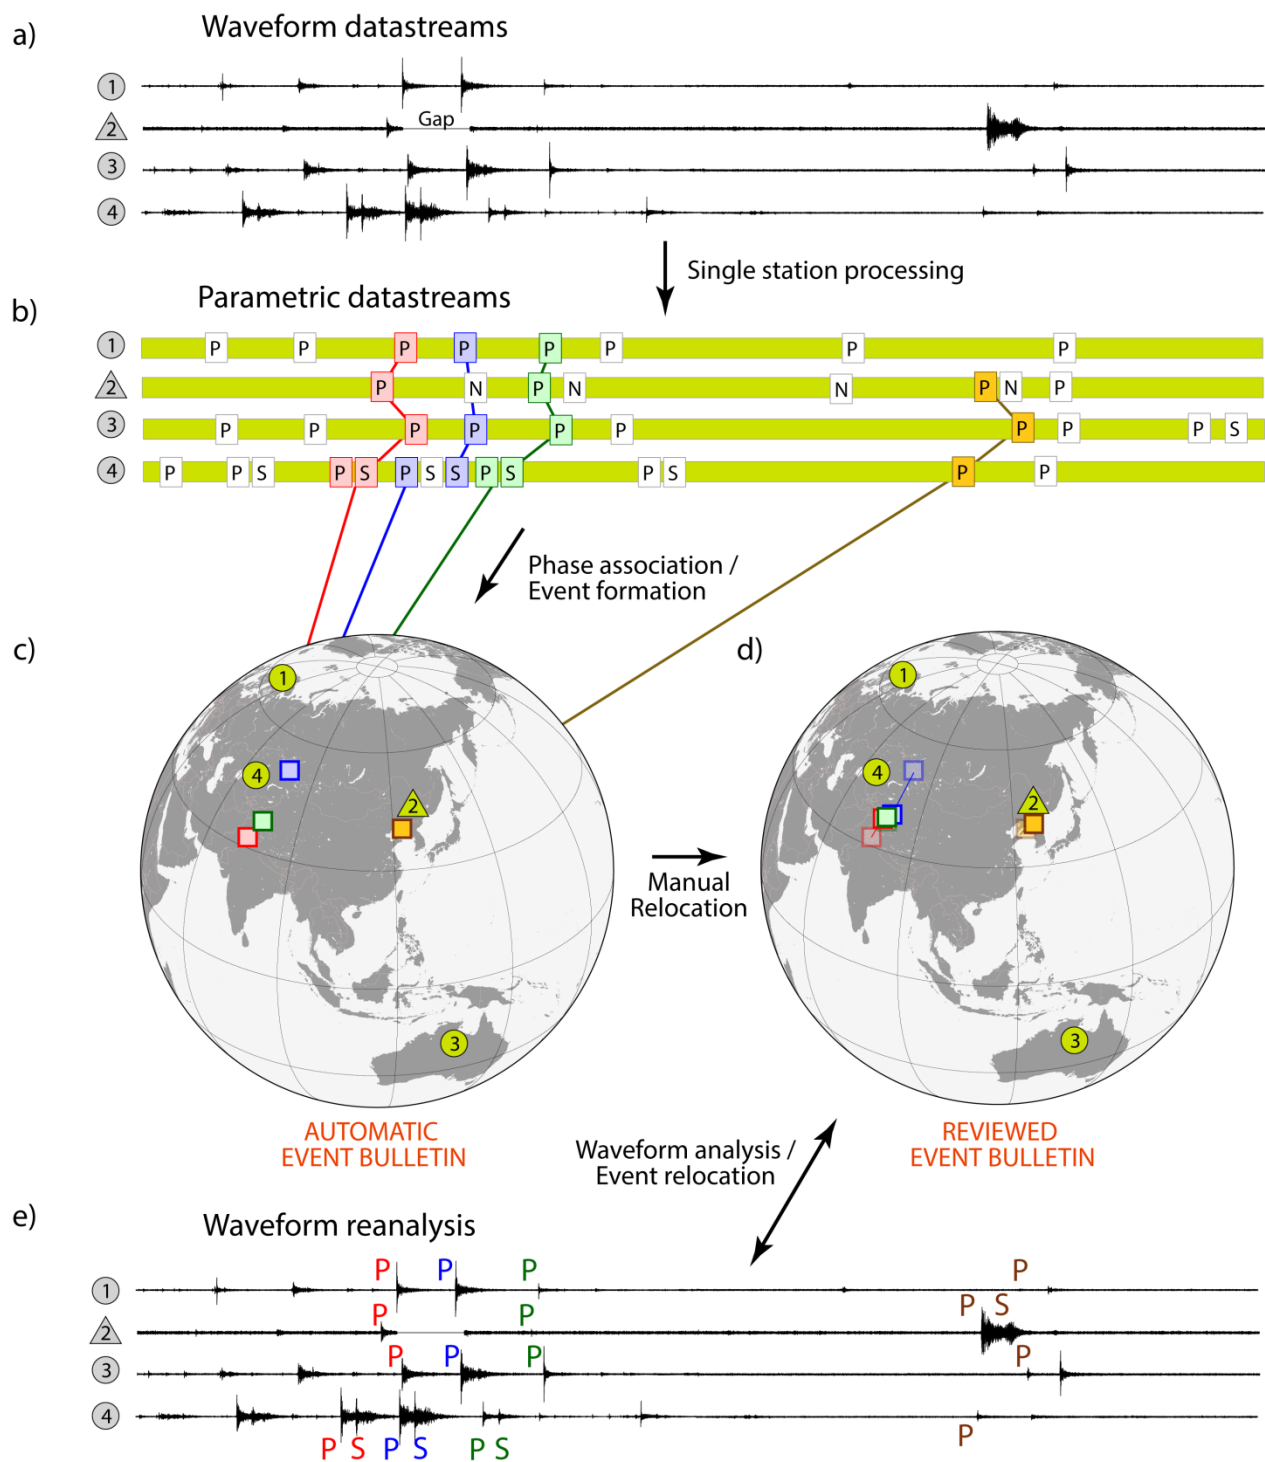


Figure 48. The classical processing pipeline displayed schematically for the case of an aftershock sequence recorded on a global network. See text for details.

We consider a situation in which it is understood that an aftershock sequence is developing. We define a geographic region of interest in which events are anticipated and have the aim of extracting as much information as possible about the evolving sequence in this target region. We could then strip the parametric detection lists of phases associated with well-defined aftershocks, therefore reducing the likelihood of false associations. In aftershock sequences from moderately sized earthquakes (e.g. up to magnitude 6.5), correlation and subspace detectors may perform exceptionally well in characterizing events with a relatively limited number of templates covering the entire source region effectively (e.g. Harris and Dodge, 2011; Junek et al., 2015). For larger events, such as the M=7.8 Nepal earthquake, the source regions are many times larger with apertures up to and exceeding 100 km. The waveform similarity between one event and the next decreases both with differences in source mechanism and with increasing hypocentral distances and it has been demonstrated for several large aftershock sequences that correlation methods are relatively ineffective at characterizing the seismicity, with only a very small proportion of detected events showing significant signal similarity to other events (Slinkard et al., 2013).

Figure 49 illustrates how we may attempt to construct a timeline for event hypotheses within the geographical region of the ongoing aftershock sequences using some form of signal processing optimized for that particular source region on the available stations (a). It is not specified in the figure which algorithm we employ for this purpose – we could be using correlators or subspace detectors (Harris and Dodge, 2011) or some form of grid search method. With an accurate list of hypocenters and origin times, we may predict with considerable accuracy the sequence of phases anticipated at the different stations and these times can be used to mask out the corresponding phases (detected or not) from the parametric data streams (b). The reduced parametric data streams can then be sent again to the phase association process (c.f. Figure 48 b and c). While we may not guarantee a more accurate association of the remaining detections, we are likely to have reduced significantly the number of false event hypotheses involving significant events from the aftershock region. We will later consider procedures for evaluating offline the validity of event hypotheses formed and an iterative process could be conceived whereby the parametric datastream is reduced repeatedly until no further changes can be achieved.

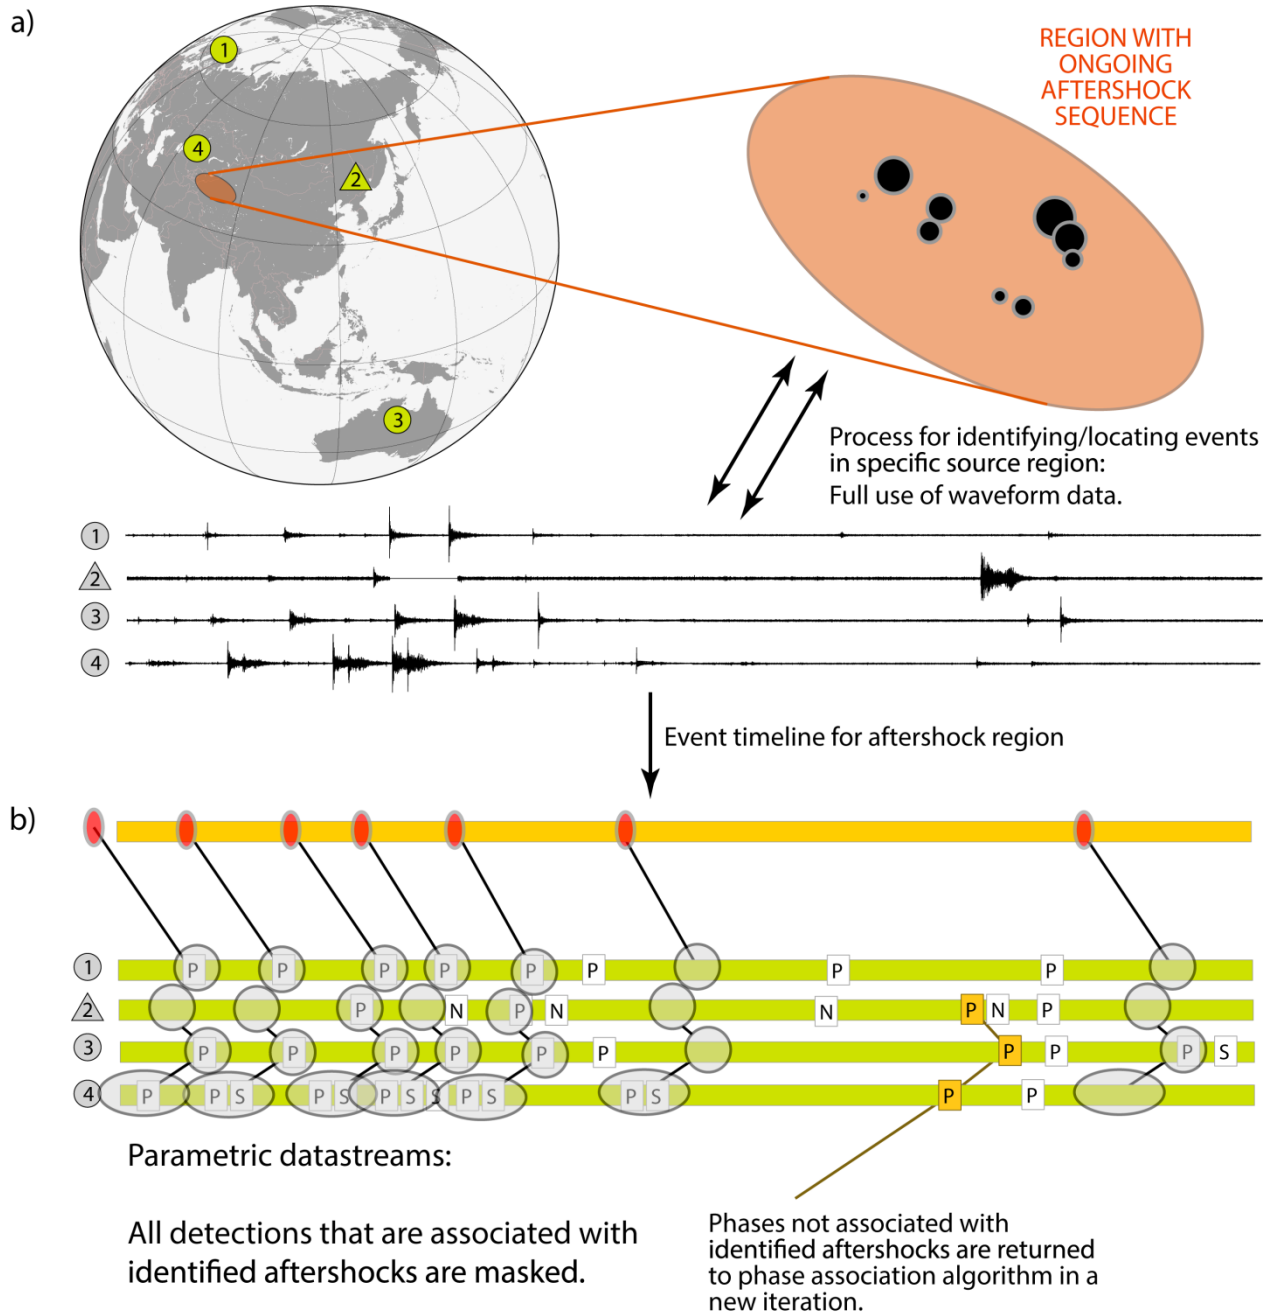


Figure 49. The scenario whereby the processing pipeline understands that an aftershock sequence is developing in a limited geographical region and spawns an offline process to detect and locate aftershocks accurately such that the detection parameters associated with well-characterized events can be removed from the parametric data stream provided to the phase association algorithm.

2.6 An Autoregressive and Grid Search Approach to Aftershock Classification

When it becomes clear that an aftershock sequence is underway, we should be able to evaluate rapidly which stations are likely to provide optimal detection capability for the given source

region (e.g. Kværna and Ringdal, 2013). New array beams, additional to those present in the initial detection beam deployment, may be formed. The uppermost trace in Figure 50 displays a beam on the ARCES array optimized for detecting teleseismic P-phases from the Nepal region (ARCES being determined to be one of the most sensitive of the IMS seismic arrays for signals from this part of the world). The lowermost trace in Figure 50 is a continuous AR-AIC function, synthesized as a summation of AR-AIC traces calculated in relatively short overlapping windows (see e.g. Leonard and Kennett, 1999) for which signal onsets are readily identified as significant local maxima. The AR-AIC trace makes it easier to observe the sequence of arrivals given the enormous dynamic range of the waveforms but, more significantly, provides far more accurate estimates of the signal onset times than is possible using trigger times from classical STA:LTA traces.




Figure 50. ARCES array beam aimed at teleseismic P-arrivals from the Nepal earthquake, together with the continuous AR-AIC function for this particular beam.

If the procedure displayed in Figure 50 is carried out on each of several of the most sensitive stations for the source region in question, with as great an azimuthal coverage as possible, then we can form preliminary event hypotheses in this source region alone, simply by finding local minima in the AR-AIC traces that propagate back to a common origin time. The most effective way of doing this in practice is to pre-calculate P-wave traveltimes from each point on the geographical grid covering the aftershock region to each of the stations selected. Given that this calculation is only performed once for the entire sequence, the best results are anticipated

when using a fully 3-dimensional P-wave velocity model (e.g. Simmons et al., 2012). An empirical threshold on the AR-AIC traces is set and every local AR-AIC maximum which exceeds this threshold is selected as a trigger. For a given trigger, we loop around every node on the geographical grid and define the origin time for a hypothetical event at that grid node. Given this candidate hypocenter, we then calculate the anticipated arrival times at each of the observing stations and seek local maxima in the AR-AIC traces close to this time. For the majority of the trial hypocenters, we will only find a significant local maximum at the time of the trigger detection. Trial hypocenters for which several local maxima fall very close to the predicted arrival times are considered to be likely candidates and are stored. The candidate hypocenter for which the greatest number of arrivals match the local maxima of the AR-AIC traces is selected as a preliminary event hypothesis. If many candidate event hypotheses predict the same number of arrivals, then the location with the smallest traveltime residual is selected. Once an event hypothesis is selected, all of the corresponding local maxima in the AR-AIC traces are set to zero and the procedure is repeated until all of the event hypotheses which are supported by a minimum number of observations (4 for example) are exhausted. While this method operates with a certain degree of brute force, the reliability of the arrival time estimates in the AR-AIC traces, coupled with the fact that all traveltimes are pre-calculated, means that the procedure runs very rapidly when only a modest number of observing stations are used.

If possible, local maxima in the AR-AIC traces that are not related to the aftershock sequence can be screened out using, for example, f-k analysis. This will reduce the number of triggers to loop around and will reduce the risk of generating spurious event hypotheses. We do not of course eliminate the risk of generating false event hypotheses, and neither do we eliminate the risk of corrupting otherwise valid event hypotheses. However, we appear to reduce the likelihood that high SNR phase arrivals corresponding to events in the aftershock sequence are associated spuriously to form event hypotheses far from the aftershock source region. Figure 51 displays the fixed depth hypocenters obtained over the first 3 days of the Nepal sequence using this method. There are 180 preliminary event hypotheses displayed here, each without a magnitude estimate. In comparison, 345 events are found in the Reviewed Event Bulletin of the IDC in this region and over the same time interval.

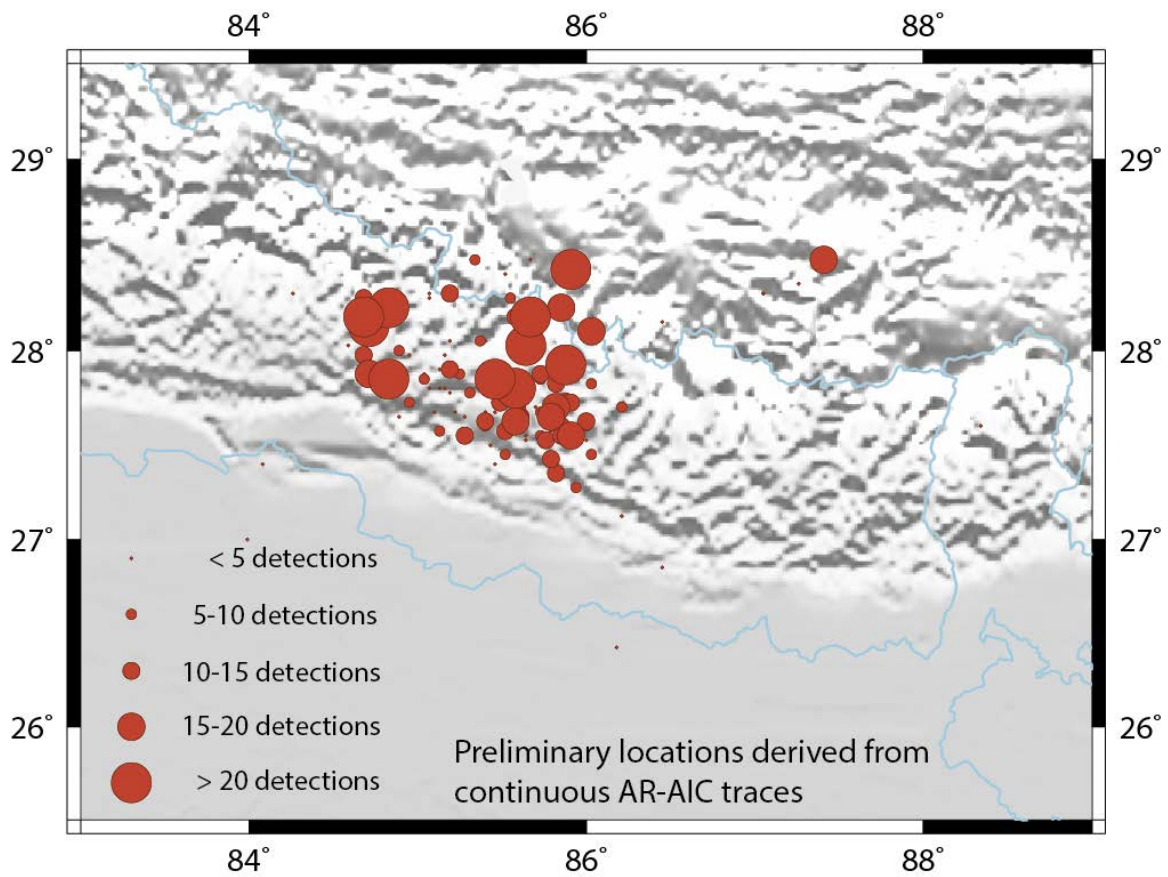


Figure 51. Fully automatic event location estimates obtained in the first 3 days of the Nepal sequence using an iterative grid search method based upon local maxima of the AR-AIC functions and pre-calculated traveltimes. Numerous event hypotheses are formed and, on each iteration, those hypotheses which best satisfy an optimum number of phase arrivals on different stations are accepted. 180 event hypotheses are displayed here, compared with 345 events in the REB in the same time period in the same region. Note that the size of the symbols relates only to the number of local maxima on the AR-AIC traces are included and is not necessarily indicative of the event magnitude.

Figure 52 displays detection lists from four of the selected arrays. In the lowermost of the two traces for each station, a vertical line is drawn for each detection for which the indicated conditions on backazimuth and apparent velocity are met and where the height of the line is proportional to the signal to noise ratio. In the uppermost traces, we display the set of detections where each arrival associated with one of the event hypotheses in Figure 3 has been removed. The number of arrivals satisfying the stated conditions has been approximately halved for each station. Significantly, all of the detections with high signal to noise ratio have been associated with one of our event hypotheses. This source region-specific feature would then either repeat with a lower detection threshold, or would simply send the unassociated detections back to the Global Association algorithm for a new iteration. This demonstrates a proof of concept of a context-based procedure for improving the processing pipeline in the large earthquake aftershock scenario.

The Nepal example illustrated here includes only teleseismic P phases since no open data from near regional distances was available for this sequence. In many circumstances, regional phases may be available and, especially within regions with dense networks, are likely to dominate the data streams. Figure 53 displays waveforms from 3 stations within 10 degrees of the magnitude 7.6 Kashmir earthquake of October 8, 2005. For all stations, one trace is optimized for the regional P phase and one for the regional S phase. (For the NIL station at a distance of approximately 1 degree, this is likely to be a crustal Pg/Sg-type phase, whereas for the more distant DEBE and KKAR stations, we are looking at Pn/Sn.) Most of the regional arrivals are visible both on traces optimized for P and traces optimized for S.

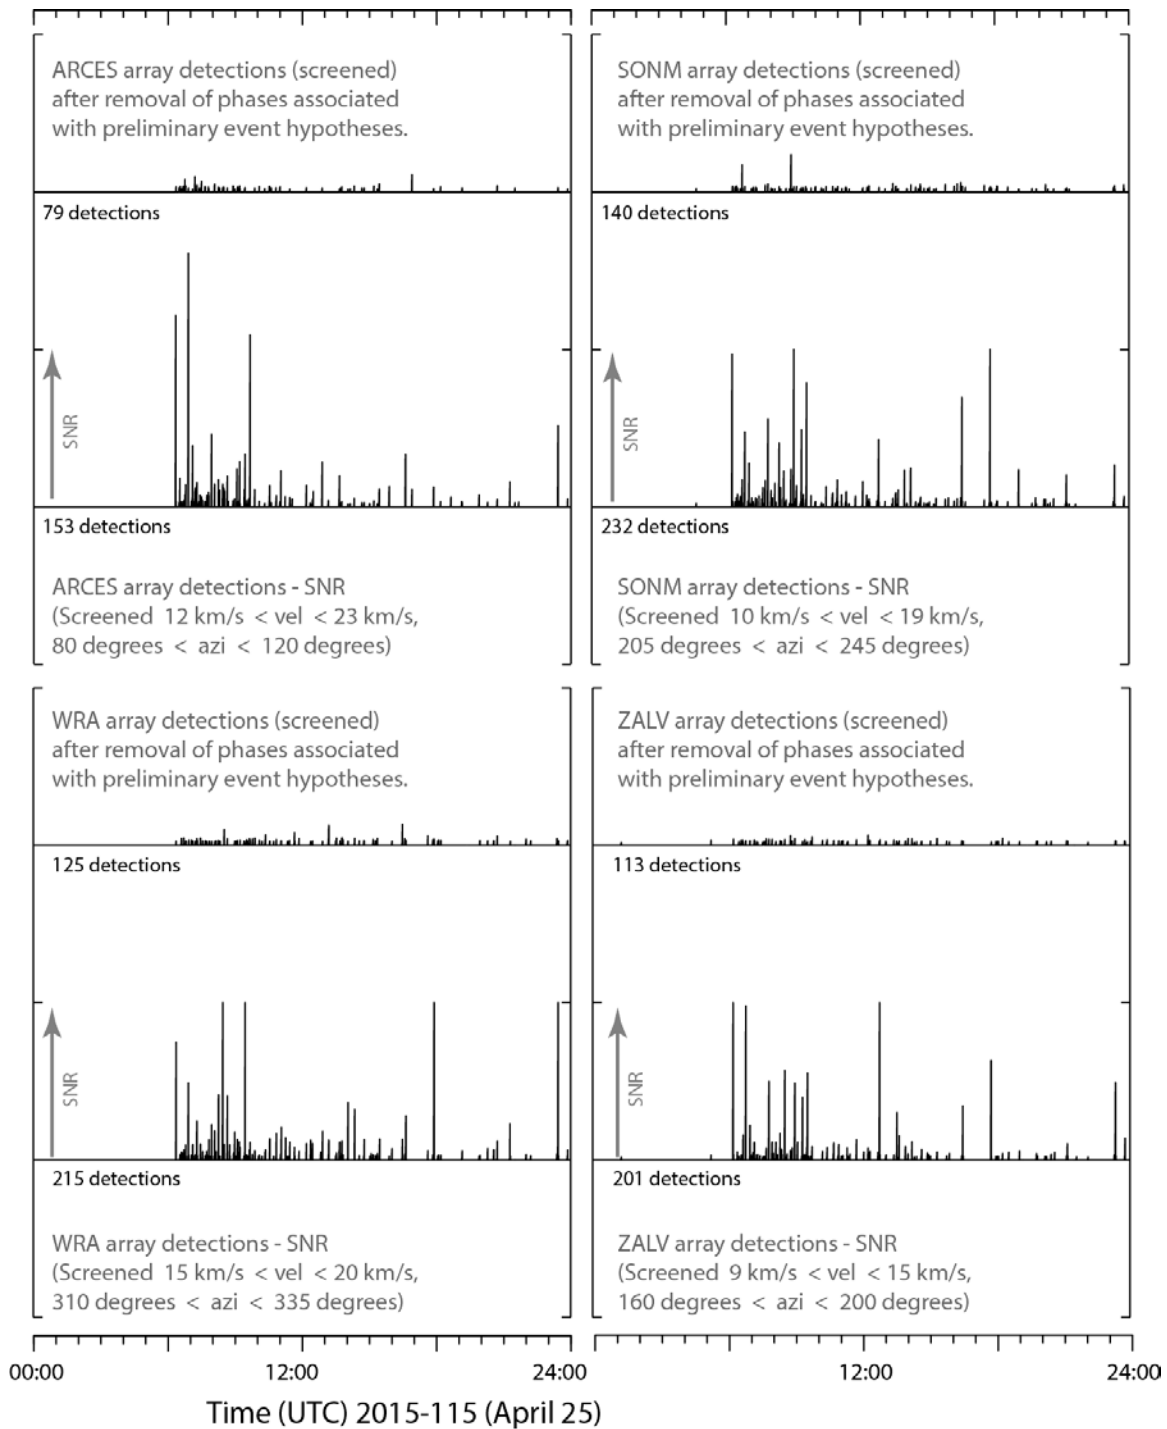


Figure 52. For each of the four array stations listed, the lowermost trace displays detections (SNR as a function of time) for which the velocity and azimuth constraints hold. The traces above display the detections remaining after phases that can be associated with the preliminary events displayed in Figure 3 are removed. All detections with the highest SNR measurements are removed from the input. The number of detections for each station is approximately halved.




Figure 53. Traces optimized for the detection of regional P and regional S at 3 stations within 10 degrees of the October 8, 2005, Kashmir earthquake. For the KKAR array, both of these traces are beams steered with parameters chosen to optimize the SNR for the sought-after phases. For the 3-component stations NIL and DEBE, we use the vertical component only for the P-phases and the transverse horizontal rotation for the S-phases. The frequency bands used are generally higher for the P-phases than for the S-phases. Note that the NIL waveforms are clipped for the largest events meaning that only the P-arrival is usually seen for these events.

Figure 54 displays the AR-AIC functions for the optimal waveforms displayed in Figure 53. As with the teleseismic case displayed in Figure 50, we see that the enormous dynamic range of the waveform amplitudes is reduced considerably in the AR-AIC traces, with the significance of the arrivals from the smaller amplitude events far more comparable to that of the much larger events.




Figure 54. Continuous AR-AIC functions for the traces displayed in Figure 53.

If we look more closely at the AIC traces in Figure 54 we see a significant peak for both P and S phases for each arrival (Figure 55). However, for every identified P phase, the peak value of the AR-AIC trace for the P-optimized waveform is almost guaranteed to exceed the peak value of the AR-AIC trace for the S-optimized waveform for that time. Similarly, for every identified S phase, the peak value of the AR-AIC trace for the S-optimized waveform is almost guaranteed to exceed the corresponding peak value of the AR-AIC trace for the P-optimized waveform. Figure 55 provides optimism that a relatively simple detection reduction scheme, where the AR-AIC trace is set to zero for the time surrounding any trigger at that station for all AR-AIC traces that do not provide the maximum value, is likely to identify most phases correctly and therefore reduce the likelihood of spurious event hypotheses.

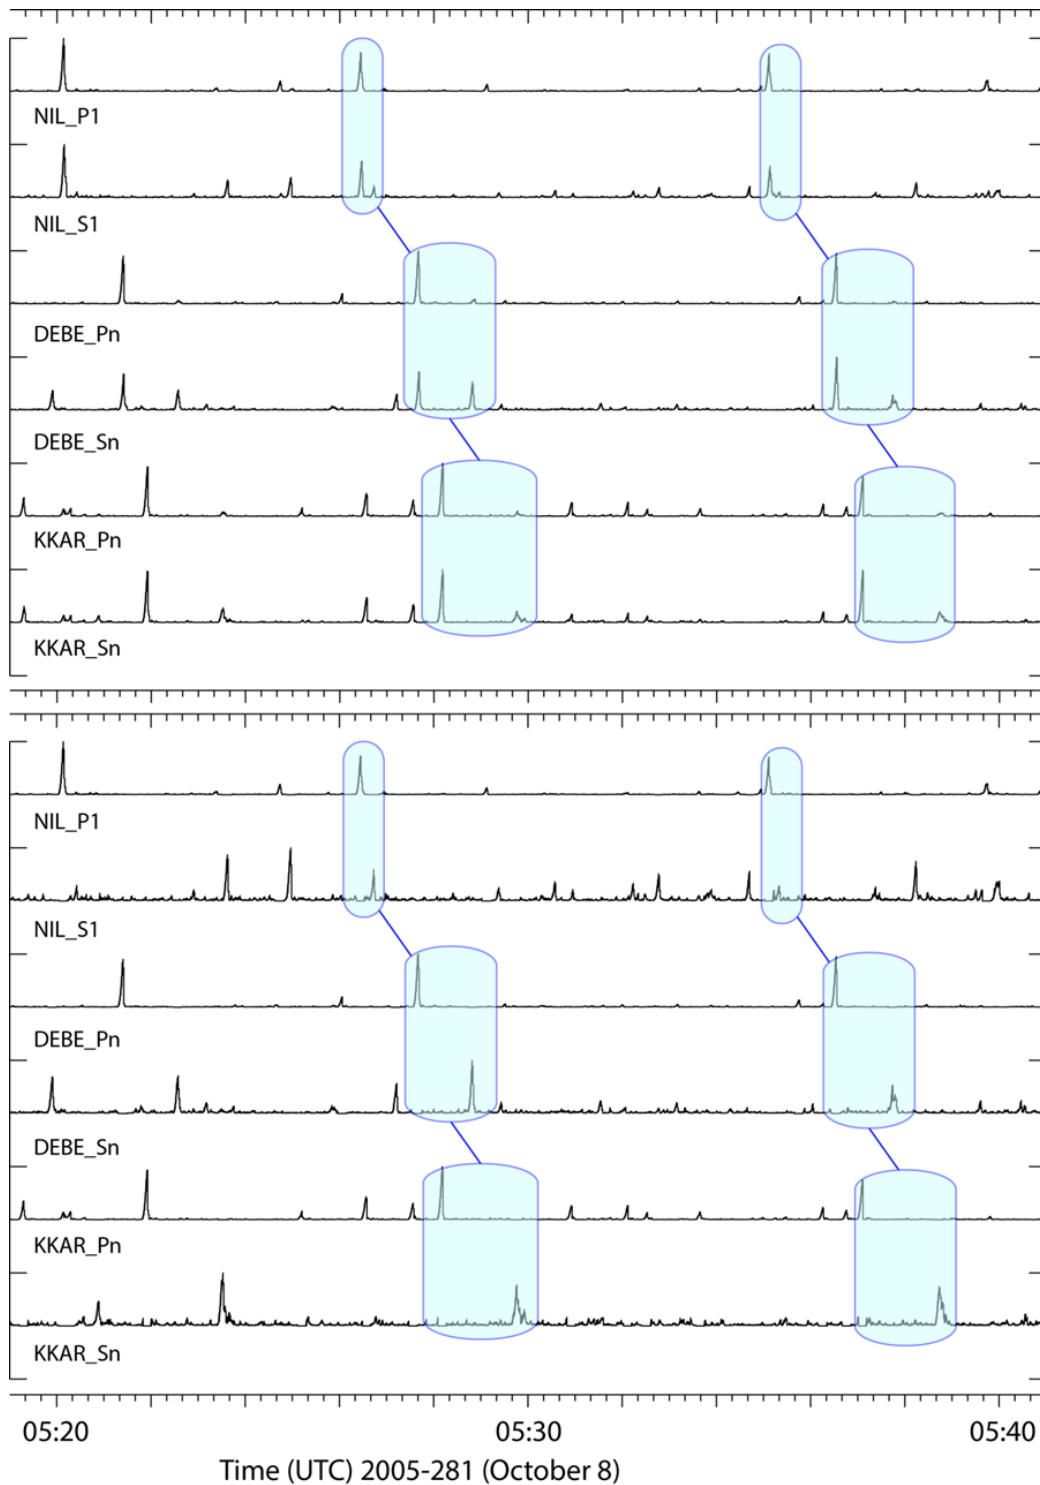


Figure 55. Close-up of the traces displayed in Figure 54 in original form (upper panel) and in reduced form (lower panel).

3 Results and Discussions

In this section we provide a brief overview of the results from each of the studies described in the previous section.

3.1 Cancellation of Repetitive Transients in Seismic Array Data Streams

We have highlighted the problem caused by repeating high amplitude transient signals from very local sources on key monitoring stations. These signals are usually genuine responses to ground motion (as opposed to artefacts in the data) but caused by sources that are not of monitoring interest. Examples are ice-quake signals generated in the glaciers surrounding the SPITS array on Svalbard and Rg-type phases generated by industrial events in the vicinity of several other stations.

Two types of algorithm are proposed: an intermittent-source algorithm and a continuous-source algorithm. The intermittent-source algorithm fails in some cases where a complicated source-time function results in the superposition of nuisance signals. The continuous-source cancellor works well in these situations. However, for practical implementation, the intermittent-source algorithm is currently recommended given the tradeoff between effectiveness and speed.

3.2 A Revised Detection Framework Architecture

The intermittent-source cancellation algorithm has been implemented into the so-called Detection Framework. In order to facilitate the implementation of new methods and outscaling data processing to a cluster of computers, the Detection Framework was successfully upgraded with a new architecture.

The performance of the revised Detection Framework with respect to cancellation of local icequake signals was evaluated using data from the SPITS array. For a three day period the cancellation provided a factor 10 reduction in nuisance icequake signals. The data “cleaned” from nuisance signals again provide for improvements in both manual and automatic data analysis.

We conducted an experiment with automatic reprocessing of the “cleaned” SPITS data within the revised Detection framework. This demonstrated improvements in detection performance, but also revealed some issues to be resolved regarding partial-only signal cancellation of time-overlapping nuisance signals.

3.3 Empirical Matched Field Processing and Adaptive Beamforming

We sought to improve the sensitivity of signal detectors using empirical matched field processing (EMFP) and adaptive beamforming. Matched Field Processing can be considered to be a form of frequency domain beamforming and EMFP is a calibrated form where the applied phase shifts are calculated by the observations of previous signals from the direction of interest. EMFP provides the optimal estimator of the signal under the assumption of white background noise. In the case of correlated noise, the optimal choice for steering vectors is the adaptive beamformer weighting – or minimum variance distortionless receiver (MVDR).

To demonstrate, signals from confirmed nuclear tests in the DPRK were submerged into aftershock sequences recorded on medium aperture seismic arrays (with aperture of the order 10 km) with the noise covariance structure for MVDR estimated from a time average of the ongoing aftershock sequence. Relative to the standard matched field estimator, adaptive processors reduce, often dramatically, the contribution of interfering events to detection statistics traces targeted on the DPRK test.

We draw two conclusions from the examples provided: (1) the use of calibrated steering vectors with empirical matched field processing appears to allow adaptive beamforming techniques to be used successfully, and (2) adaptive beamforming may be used to advantage in several stages of iteration in an advanced pipeline. It may support more sensitive initial event detection in the pipeline. In addition, adaptive beamforming may support context-driven reprocessing of a stream at a station which failed to make a detection when observations at other stations in a network indicate that it should have had a trigger.

3.4 Evaluation of Event Hypotheses

Fully automatic phase association algorithms form the basis for the seismic event bulletins generated by data centers. Phase detections from single stations are grouped together from the parametric datastreams by algorithms such as GA (Global Association) or NetVISA. Under current operations, only at the analyst review stage is it possible to evaluate these hypotheses and revise them.

In an iterative processing pipeline, we should have the means by which to evaluate event hypotheses automatically using processes operating on the waveform data. Using historical observations, recorded on permanent networks, it is usually possible to obtain a reasonable estimate of the detectability for an event of a given size, in a given location, and recorded on a given station. The absence of evidence of a signal at a time at which an arrival is predicted from an event hypothesis should alert an automatic procedure to a likely incorrect hypothesis once explanations such as data dropouts or abnormal background noise have been eliminated.

We discuss the limitations of fully automatic hypothesis evaluation with a test case of two events in Central Asia – a deep Hindu Kush earthquake and a shallow earthquake in Kashmir approximately a minute later. Almost all of the arrivals from the Kashmir event are buried in the coda of the Hindu Kush event. Automatic processes can readily identify the false automatic event hypothesis which includes phases from the Kashmir event but would struggle to identify the buried arrivals with sufficiently high confidence. It is exactly such a scenario for which the adaptive beamforming/MVDR process described in the previous section is envisaged.

3.5 Iterative Schemes for Automatic Event Bulletins: Coping with Aftershock Sequences

Aftershock sequences following very large earthquakes present enormous challenges to near-realtime generation of seismic bulletins. The increase in analyst resources needed to relocate an inflated number of events is compounded by failures of phase association algorithms and a significant deterioration in the quality of underlying fully automatic event bulletins. We consider the scenario where a large earthquake has occurred and propose to define a region of likely aftershock activity in which events are detected and accurately located using a separate specially targeted semi-automatic process. This effort may focus on so-called pattern detectors, but here we demonstrate a more general grid search algorithm which may cover wider source regions without requiring waveform similarity. Given many well-located aftershocks within our source region, we may remove all associated phases from the original detection lists prior to a new iteration of the phase association algorithm.

We provide a proof-of-concept example for the 2015 Gorkha sequence, Nepal, recorded on seismic arrays of the International Monitoring System. Even with very conservative conditions for defining event hypotheses within the aftershock source region, we can automatically remove about half of the original detections which could have been generated by Nepal earthquakes and reduce the likelihood of false associations and spurious event hypotheses. Further reductions in the number of detections in the parametric data streams are likely using correlation and subspace detectors and/or empirical matched field processing.

4 Conclusions

The objectives of this study were to explore some of the opportunities for more sophisticated detection algorithms we believe are feasible in an architecture allowing reprocessing of the seismic data stream. In existing processing pipelines, the stages of the automatic processing pipelines which consider the waveform data perform only a single pass of the waveforms before supplying later stages of the pipeline with parametric data from which hypotheses of seismic events are generated. Only at the final stage, where a human analyst considers the list of event hypotheses, is it currently possible for a re-evaluation of the incoming waveform data stream. We believe that a much more accurate and realistic interpretation of the incoming data can be made far earlier, in the automatic processing, with contextual signal processing performed iteratively to optimize detection and characterization of seismic events.

In this study we contemplate future pipelines that, in addition to implementing contextual signal processing, could profit from a geographically hierarchical structure organized to have a clean distributed implementation exploiting data locality. We imagine at least two levels to the hierarchy: a base level of collections of sub-pipelines processing data from individual stations, and a top level global pipeline integrating results from the base level. Whether a third, intermediate level processing regional station aggregations would add significantly to overall effectiveness has not been addressed. The value of the local level is that adaptive, contextual processing to reject local interference can take place independent of other stations, without the overall system needing to form hypotheses about or conduct processing to suppress interference observed only at one station. This is not to say that the global level of the hierarchy would not provide contextual hints to local sub-pipelines regarding events or interference expected to be observed widely.

At the single station level, detection and estimation capability may be compromised significantly by the occurrence of high amplitude transients from local sources. We prototype an architectural element to detect, and subtract from the stream, local nuisance transient signals that plague critical monitoring arrays. In addition, we test means to suppress aftershocks with adaptive beamforming algorithms given information about numerous signals present in the current analysis window, for example from ongoing aftershock sequences. This is an example where global context would inform local signal processing adaptations.

At the global level, we consider how recognition of ongoing aftershock sequences could affect contextual event association. Aftershock sequences following very large earthquakes present enormous challenges to near-real-time generation of seismic bulletins both due to the increased

numbers of events and a decrease in the quality of automatic bulletins. We consider the scenario where a large earthquake has occurred and propose to define a region of likely aftershock activity in which events are detected and accurately located using a separate specially targeted semi-automatic process. Given many well-located aftershocks within our source region, we may remove all associated phases from the original detection lists prior to a new iteration of the phase association algorithm. We provide a proof-of-concept example for the 2015 Gorkha sequence, Nepal, recorded on seismic arrays of the International Monitoring System. Even with very conservative conditions for defining event hypotheses within the aftershock source region, we can automatically remove about half of the original detections which could have been generated by Nepal earthquakes and reduce the likelihood of false associations and spurious event hypotheses. Further reductions in the number of detections in the parametric data streams are likely using correlation and subspace detectors and/or empirical matched field processing.

References

- Arora, N. S., Russell, S., and Sudderth, E. (2013), NET-VISA: Network Processing Vertically Integrated Seismic Analysis, *Bulletin of the Seismological Society of America*, 103, pp. 709-729, doi:10.1785/0120120107.
- Capon, J., Greenfield, R. J., and Kolker, R. J. (1967), Multidimensional maximum-likelihood processing of a large aperture seismic array, *Proc. IEEE*, 55 (2), pp. 192-211.
- Gibbons, S. J., Schweitzer, J., Ringdal, F., Kværna, T., Mykkeltveit, S., and B. Paulsen (2011), Improvements to Seismic Monitoring of the European Arctic Using Three-Component Array Processing at SPITS, *Bulletin of the Seismological Society of America*, Vol. 101, No. 6, pp. 2737-2754, doi:10.1785/0120110109.
- Harris, D. B. (1989), Characterizing source regions with signal subspace methods: theory and computational methods, UCID-21848, Lawrence Livermore National Laboratory, Livermore, CA.
- Harris, D. B. (2006), Subspace Detectors: Theory, UCRL-TR-222758, Lawrence Livermore National Laboratory, Livermore, CA.
- Harris, D. B., Jarpe, S. P., and Harben, P. E. (1991), Seismic noise cancellation in a geothermal field, *Geophysics*, 56(10), pp. 1677-1680.
- Harris, D. B. and Dodge, D. A. (2011), An Autonomous System for Grouping Events in a Developing Aftershock Sequence, *Bulletin of the Seismological Society of America*, Vol. 101, No. 2 (1 April 2011), pp. 763-774, doi:10.1785/0120100103.
- Junek, W. N., Kværna, T., Pirli, M., Schweitzer, J., Harris, D. B., Dodge, D. A., and Woods, M. T. (2015), Inferring Aftershock Sequence Properties and Tectonic Structure Using Empirical Signal Detectors, *Pure and Applied Geophysics*, 172, pp. 359-373, doi:10.1007/s00024-014-0938-0.
- Kværna, T. and Ringdal, F. (2013), Detection Capability of the Seismic Network of the International Monitoring System for the Comprehensive Nuclear-Test-Ban Treaty, *Bulletin of the Seismological Society of America*, 103, pp. 759-772, doi:10.1785/0120120248.
- Kværna, T., Harris, D. B., Gibbons, S. J., and Dodge, D. A. (2014), Adapting Pipeline Architectures to Track Developing Aftershock Sequences and Recurrent Explosions, AFRL-RV-PS-TR-2014-0075, NORSAR, Kjeller, Norway.
- Leonard, M., and Kennett, B. L. N. (1999), Multi-component autoregressive techniques for the analysis of seismograms, *Physics of the Earth and Planetary Interiors*, 113, pp. 247-263, doi:10.1016/s0031-9201(99)00054-0.
- Murphy, J. R. and Barker, B. W. (2003), Revised Distance and Depth Corrections for Use in the Estimation of Short-Period P-Wave Magnitudes, *Bull. Seism. Soc. Am.*, Vol. 93 (4), pp. 1746-1764, doi:10.1785/0120020084.
- Oppenheim, A. V. and Schafer, R. W. (1975), *Digital Signal Processing*, Englewood Cliffs, N.J., Prentice-Hall. ISBN 0-13-214635-5.

- Simmons, N. A., Myers, S. C., Johannesson, G., and Matzel, E. (2012), LLNL-G3Dv3: Global P wave tomography model for improved regional and teleseismic travel time prediction, *Journal of Geophysical Research*, 117, No. B10., B10302, doi:10.1029/2012jb009525.
- Slinkard, M. E., Carr, D. B., and Young, C. J. (2013), Applying Waveform Correlation to Three Aftershock Sequences, *Bulletin of the Seismological Society of America*, Vol. 103, No. 2A (01 April 2013), pp. 675-693, doi:10.1785/0120120058.
- Widrow, B. and Hoff, M. E. (1960), Adaptive switching circuits, *Proc. Of WESCON Conv. Rec.*, Part 4, pp. 96-140.

Appendix: Levinson-Durbin Algorithm for Solution of Block Toeplitz Matrix Equation

The system of equations requiring solution involves a block Toeplitz matrix:

$$\begin{bmatrix} \mathbf{R}_0 & \mathbf{R}_1 & \cdots & \mathbf{R}_k \\ \mathbf{R}_1^T & \mathbf{R}_0 & \mathbf{R}_1 & \vdots \\ \vdots & \mathbf{R}_1^T & \ddots & \mathbf{R}_1 \\ \mathbf{R}_k^T & \cdots & \mathbf{R}_1^T & \mathbf{R}_0 \end{bmatrix} \begin{bmatrix} \mathbf{a}[0] \\ \mathbf{a}[1] \\ \vdots \\ \mathbf{a}[k] \end{bmatrix} = \begin{bmatrix} \mathbf{C}[0] \\ \mathbf{C}[1] \\ \vdots \\ \mathbf{C}[k] \end{bmatrix} \quad (\text{A-1})$$

The solution proceeds by recursion. Suppose we have the solution for the $(k-1)^{\text{st}}$ order problem:

$$\begin{bmatrix} \mathbf{R}_0 & \mathbf{R}_1 & \cdots & \mathbf{R}_{k-1} \\ \mathbf{R}_1^T & \mathbf{R}_0 & \mathbf{R}_1 & \vdots \\ \vdots & \mathbf{R}_1^T & \ddots & \mathbf{R}_1 \\ \mathbf{R}_{k-1}^T & \cdots & \mathbf{R}_1^T & \mathbf{R}_0 \end{bmatrix} \begin{bmatrix} \mathbf{a}^{k-1}[0] \\ \mathbf{a}^{k-1}[1] \\ \vdots \\ \mathbf{a}^{k-1}[k-1] \end{bmatrix} = \begin{bmatrix} \mathbf{C}[0] \\ \mathbf{C}[1] \\ \vdots \\ \mathbf{C}[k-1] \end{bmatrix} \quad (\text{A-2})$$

We construct the solution to the k^{th} order problem given that of the $(k-1)^{\text{st}}$ order problem. Note that:

$$\begin{bmatrix} & & & \mathbf{R}_k \\ & [\mathbf{R}]_{k-1} & & \mathbf{R}_{k-1} \\ & & \ddots & \vdots \\ \mathbf{R}_k^T & \mathbf{R}_{k-1}^T & \cdots & \mathbf{R}_0 \end{bmatrix} \begin{bmatrix} \mathbf{a}^k[0] \\ \mathbf{a}^k[1] \\ \vdots \\ \mathbf{a}^k[k] \end{bmatrix} = \begin{bmatrix} \mathbf{C}[0] \\ \mathbf{C}[1] \\ \vdots \\ \mathbf{C}[k] \end{bmatrix} \quad (\text{A-3})$$

and:

$$\begin{bmatrix} \mathbf{R}_0 & \cdots & \mathbf{R}_{k-1} & \mathbf{R}_k \\ \vdots & & & \\ \mathbf{R}_{k-1}^T & & [\mathbf{R}]_{k-1} & \\ \mathbf{R}_k^T & & & \end{bmatrix} \begin{bmatrix} \mathbf{a}^k[0] \\ \mathbf{a}^k[1] \\ \vdots \\ \mathbf{a}^k[k] \end{bmatrix} = \begin{bmatrix} \mathbf{C}[0] \\ \mathbf{C}[1] \\ \vdots \\ \mathbf{C}[k] \end{bmatrix} \quad (\text{A-4})$$

where:

$$[\mathbf{R}]_{k-1} = \begin{bmatrix} \mathbf{R}_0 & \mathbf{R}_1 & \cdots & \mathbf{R}_{k-1} \\ \mathbf{R}_1^T & \mathbf{R}_0 & \mathbf{R}_1 & \vdots \\ \vdots & \mathbf{R}_1^T & \ddots & \mathbf{R}_1 \\ \mathbf{R}_{k-1}^T & \cdots & \mathbf{R}_1^T & \mathbf{R}_0 \end{bmatrix} \quad (\text{A-5})$$

The solution proceeds by carrying along recursions for two special cases in addition to the main problem of (A.1):

$$\begin{bmatrix} \mathbf{R}_0 & \mathbf{R}_1 & \cdots & \mathbf{R}_k \\ \mathbf{R}_1^T & \mathbf{R}_0 & \mathbf{R}_1 & \vdots \\ \vdots & \mathbf{R}_1^T & \ddots & \mathbf{R}_1 \\ \mathbf{R}_k^T & \cdots & \mathbf{R}_1^T & \mathbf{R}_0 \end{bmatrix} \begin{bmatrix} \alpha^k[0] \\ \alpha^k[1] \\ \vdots \\ \alpha^k[k] \end{bmatrix} = \begin{bmatrix} \mathbf{I} \\ \mathbf{0} \\ \vdots \\ \mathbf{0} \end{bmatrix} \quad (\text{A-6})$$

and:

$$\begin{bmatrix} \mathbf{R}_0 & \mathbf{R}_1 & \cdots & \mathbf{R}_k \\ \mathbf{R}_1^T & \mathbf{R}_0 & \mathbf{R}_1 & \vdots \\ \vdots & \mathbf{R}_1^T & \ddots & \mathbf{R}_1 \\ \mathbf{R}_k^T & \cdots & \mathbf{R}_1^T & \mathbf{R}_0 \end{bmatrix} \begin{bmatrix} \beta^k[0] \\ \beta^k[1] \\ \vdots \\ \beta^k[k] \end{bmatrix} = \begin{bmatrix} \mathbf{0} \\ \mathbf{0} \\ \vdots \\ \mathbf{I} \end{bmatrix} \quad (\text{A-7})$$

By inspection, from (54) and (57) we have:

$$\begin{bmatrix} \mathbf{R}_0 & \mathbf{R}_1 & \cdots & \mathbf{R}_k \\ \mathbf{R}_1^T & \mathbf{R}_0 & \mathbf{R}_1 & \vdots \\ \vdots & \mathbf{R}_1^T & \ddots & \mathbf{R}_1 \\ \mathbf{R}_k^T & \cdots & \mathbf{R}_1^T & \mathbf{R}_0 \end{bmatrix} \begin{bmatrix} \alpha^{k-1}[0] \\ \alpha^{k-1}[1] \\ \vdots \\ \mathbf{0} \end{bmatrix} = \begin{bmatrix} \mathbf{I} \\ \mathbf{0} \\ \vdots \\ \mathbf{Y}_k \end{bmatrix} \quad (\text{A-8})$$

with:

$$\mathbf{Y}_k = \sum_{i=0}^{k-1} \mathbf{R}_{k-i}^T \alpha^{k-1}[i] \quad (\text{A-9})$$

From (55) and (58), we have a similar relation:

$$\begin{bmatrix} \mathbf{R}_0 & \mathbf{R}_1 & \cdots & \mathbf{R}_k \\ \mathbf{R}_1^T & \mathbf{R}_0 & \mathbf{R}_1 & \vdots \\ \vdots & \mathbf{R}_1^T & \ddots & \mathbf{R}_1 \\ \mathbf{R}_k^T & \cdots & \mathbf{R}_1^T & \mathbf{R}_0 \end{bmatrix} \begin{bmatrix} \mathbf{0} \\ \beta^{k-1}[0] \\ \vdots \\ \beta^{k-1}[k-1] \end{bmatrix} = \begin{bmatrix} \mu_k \\ \mathbf{0} \\ \vdots \\ \mathbf{I} \end{bmatrix} \quad (\text{A-10})$$

with:

$$\boldsymbol{\mu}_k = \sum_{i=0}^{k-1} \mathbf{R}_{i+1} \boldsymbol{\beta}^{k-1}[i] \quad (\text{A-11})$$

We need to find matrices $\boldsymbol{\theta}_k$ and $\boldsymbol{\varphi}_k$ s.t.

$$\begin{bmatrix} \mathbf{I} \\ \mathbf{0} \\ \vdots \\ \mathbf{Y}_k \end{bmatrix} \boldsymbol{\theta}_k + \begin{bmatrix} \boldsymbol{\mu}_k \\ \mathbf{0} \\ \vdots \\ \mathbf{I} \end{bmatrix} \boldsymbol{\varphi}_k = \begin{bmatrix} \mathbf{I} \\ \mathbf{0} \\ \vdots \\ \mathbf{0} \end{bmatrix} \quad (\text{A-12})$$

There are two non-trivial constraints in (63):

$$\boldsymbol{\theta}_k + \boldsymbol{\mu}_k \boldsymbol{\varphi}_k = \mathbf{I} \quad (\text{A-13})$$

$$\mathbf{Y}_k \boldsymbol{\theta}_k + \boldsymbol{\varphi}_k = \mathbf{0}$$

Solving for $\boldsymbol{\theta}_k$:

$$\boldsymbol{\theta}_k - \boldsymbol{\mu}_k \mathbf{Y}_k \boldsymbol{\theta}_k = \mathbf{I} \quad (\text{A-14})$$

$$\boldsymbol{\theta}_k = (\mathbf{I} - \boldsymbol{\mu}_k \mathbf{Y}_k)^{-1} \quad (\text{A-15})$$

Solving for $\boldsymbol{\varphi}_k$:

$$\boldsymbol{\varphi}_k = -\mathbf{Y}_k (\mathbf{I} - \boldsymbol{\mu}_k \mathbf{Y}_k)^{-1} \quad (\text{A-16})$$

Thus:

$$\begin{bmatrix} \boldsymbol{\alpha}^k[0] \\ \boldsymbol{\alpha}^k[1] \\ \vdots \\ \boldsymbol{\alpha}^k[k] \end{bmatrix} = \begin{bmatrix} \boldsymbol{\alpha}^{k-1}[0] \\ \boldsymbol{\alpha}^{k-1}[1] \\ \vdots \\ \mathbf{0} \end{bmatrix} (\mathbf{I} - \boldsymbol{\mu}_k \mathbf{Y}_k)^{-1} - \begin{bmatrix} \mathbf{0} \\ \boldsymbol{\beta}^{k-1}[0] \\ \vdots \\ \boldsymbol{\beta}^{k-1}[k-1] \end{bmatrix} \mathbf{Y}_k (\mathbf{I} - \boldsymbol{\mu}_k \mathbf{Y}_k)^{-1} \quad (\text{A-17})$$

Similarly, for the $\beta^k[i]$, we need to find matrices $\hat{\theta}_k$ and $\hat{\phi}_k$ s.t.

$$\begin{bmatrix} \mathbf{I} \\ \mathbf{0} \\ \vdots \\ \mathbf{Y}_k \end{bmatrix} \hat{\theta}_k + \begin{bmatrix} \mu_k \\ \mathbf{0} \\ \vdots \\ \mathbf{I} \end{bmatrix} \hat{\phi}_k = \begin{bmatrix} \mathbf{0} \\ \mathbf{0} \\ \vdots \\ \mathbf{I} \end{bmatrix} \quad (\text{A-18})$$

Again there are two non-trivial constraints:

$$\hat{\theta}_k + \mu_k \hat{\phi}_k = \mathbf{0} \quad (\text{A-19})$$

$$\mathbf{Y}_k \hat{\theta}_k + \hat{\phi}_k = \mathbf{I}$$

With the result:

$$\hat{\phi}_k = (\mathbf{I} - \mathbf{Y}_k \mu_k)^{-1} \quad (\text{A-20})$$

$$\hat{\theta}_k = -\mu_k (\mathbf{I} - \mathbf{Y}_k \mu_k)^{-1} \quad (\text{A-21})$$

Thus:

$$\begin{bmatrix} \beta^k[0] \\ \beta^k[1] \\ \vdots \\ \beta^k[k] \end{bmatrix} = - \begin{bmatrix} \alpha^{k-1}[0] \\ \alpha^{k-1}[1] \\ \vdots \\ \mathbf{0} \end{bmatrix} \mu_k (\mathbf{I} - \mathbf{Y}_k \mu_k)^{-1} + \begin{bmatrix} \mathbf{0} \\ \beta^{k-1}[0] \\ \vdots \\ \beta^{k-1}[k-1] \end{bmatrix} (\mathbf{I} - \mathbf{Y}_k \mu_k)^{-1} \quad (\text{A-22})$$

Returning to the solution of the original problem, we try:

$$\begin{bmatrix} [\mathbf{R}]_{k-1} & \mathbf{R}_k \\ \mathbf{R}_k^T & \mathbf{R}_{k-1}^T & \cdots & \mathbf{R}_0 \end{bmatrix} \begin{bmatrix} \mathbf{a}^{k-1}[0] \\ \mathbf{a}^{k-1}[1] \\ \vdots \\ \mathbf{0} \end{bmatrix} = \begin{bmatrix} \mathbf{C}[0] \\ \mathbf{C}[1] \\ \vdots \\ \boldsymbol{\varepsilon}_k \end{bmatrix} \quad (\text{A-23})$$

where:

$$\boldsymbol{\varepsilon}_k = \sum_{i=0}^{k-1} \mathbf{R}_{k-i}^T \mathbf{a}^{k-1}[i] \quad (\text{A-24})$$

To update the $\mathbf{a}^k[i]$, we need:

$$\begin{bmatrix} \mathbf{c}[0] \\ \mathbf{c}[1] \\ \vdots \\ \boldsymbol{\varepsilon}_k \end{bmatrix} + \begin{bmatrix} \mathbf{0} \\ \mathbf{0} \\ \vdots \\ \mathbf{I} \end{bmatrix} (\mathbf{c}[k] - \boldsymbol{\varepsilon}_k) = \begin{bmatrix} \mathbf{c}[0] \\ \mathbf{c}[1] \\ \vdots \\ \mathbf{c}[k] \end{bmatrix} \quad (\text{A-25})$$

By inspection:

$$\begin{bmatrix} \mathbf{a}^k[0] \\ \mathbf{a}^k[1] \\ \vdots \\ \mathbf{a}^k[k] \end{bmatrix} = \begin{bmatrix} \mathbf{a}^{k-1}[0] \\ \mathbf{a}^{k-1}[1] \\ \vdots \\ \mathbf{0} \end{bmatrix} + \begin{bmatrix} \boldsymbol{\beta}^k[0] \\ \boldsymbol{\beta}^k[1] \\ \vdots \\ \boldsymbol{\beta}^k[k] \end{bmatrix} (\mathbf{c}[k] - \boldsymbol{\varepsilon}_k) \quad (\text{A-26})$$

Abbreviations

AFTAC	– Air Force Technical Applications Center
CTBT	– Comprehensive Nuclear-Test-Ban Treaty
CTBTO	– Comprehensive Nuclear-Test-Ban Treaty Organization
EMFP	– Empirical Matched Field Processing
IDC	– International Data Center
IMS	– International Monitoring System
ISC	– International Seismological Center
LLNL	– Lawrence Livermore National Laboratory
REB	– Reviewed Event Bulletin
SNR	– Signal-to-Noise Ratio
sps	– Samples Per Second

DISTRIBUTION LIST

DTIC/OCP 8725 John J. Kingman Rd, Suite 0944 Ft Belvoir, VA 22060-6218	1 cy
AFRL/RVIL Kirtland AFB, NM 87117-5776	2 cys
Official Record Copy AFRL/RVBYE/Robert Raistrick	1 cy

This page is intentionally left blank.